An-Nego: LLM-based Strategic Anchoring Method for Multi-Issue Negotiations

Anonymous ACL submission

Abstract

LLM-powered negotiation agents must accurately identify and respond to other participants' claims and interests to reach consensus. However, most prior work has focused on bilateral negotiation, leaving multi-party and multi-issue settings relatively underexplored, despite their prevalence in real-world semi-collaborative scenarios. We propose An-Nego, an anchor-based LLM-powered negotiation framework for multi-party stakeholder negotiation over multiple issues. An-Nego structures the process around anchored deal sets, participant voting, and moderator feedback to iteratively guide proposals and revisions. We evaluate An-Nego in terms of agreement success rate and time cost in multi-issue negotiation tasks. Experimental results show that An-Nego consistently outperforms baselines across multiple metrics, indicating improved agent capabilities in stakeholder negotiation environments. In addition, we further validate the effectiveness of the anchor effect in the multi-party gaming process. We release our code and prompts for reproducibility.[1]

1 Introduction

Recent advances in Large Language Models (LLMs) have propelled automated negotiation. Early work has primarily focused on bilateral settings (Fu et al., 2023), where two parties engage in simple deal-making. However, many real-world negotiations such as business partnerships, policy-making, and organizational coordination—involve multiple stakeholders over multiple issues, forming a complex, semi-collaborative scenario that remains underexplored (Abdelnabi et al., 2024b). In these environments, agents not only maximize their own utility but also manage conflicting interests and evolving strategies to reach consensus.

Existing LLM-based negotiation frameworks, such as LLM-Deliberation (Abdelnabi et al., 2024a), have begun to explore multi-party interactions but often lack a structured mechanism to strategically guide the negotiation process. This can lead to inefficient communication and suboptimal outcomes, particularly when dealing with multiple interrelated issues (Luo et al., 2025). Drawing from the well-established concept of the anchoring effect in behavioral economics (Tversky and Kahneman, 1974)—where an initial reference point strongly influences subsequent decisions—we posit that integrating a dynamic anchoring strategy can significantly enhance the efficiency and effectiveness of LLM-powered negotiators in complex multi-party environments.

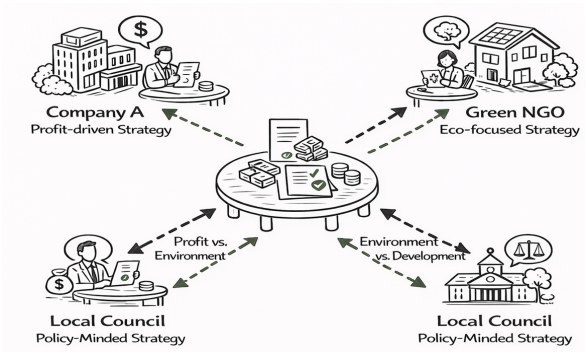

Figure 1: Multi-player and multi-issue scenarios

Our main contributions are as follows:

- We propose An-Nego, a novel LLM-powered agent framework for multi-party, multi-issue negotiations that organises interaction via an iterative *propose–vote–veto* loop.
- We introduce a strategic anchoring method where a proposer categorizes deal sets, progressively narrowing negotiations to a mutually acceptable solution space.
- Experiments on a stakeholder negotiation benchmark demonstrate that An-Nego achieves higher agreement success rates and improved dialogue quality compared to strong

[1] https://anonymous.4open.science/r/review_release-8C17

baselines.

2 Related Work

Negotiation in dialogue systems. Negotiation has long been studied as a goal-driven dialogue problem where agents must jointly optimize language and outcomes under mixed incentives. Early work in NLP explored end-to-end negotiation dialogue and planning via rollouts and reinforcement learning, showing that agents must learn both linguistic realization and strategic behavior to reach agreements (Lewis et al., 2017). Recent research further investigates how models can improve negotiation performance through self-play and feedback, indicating that iterative interaction can enhance strategy but may also increase the risk of deal failure (Fu et al., 2023).

LLM-based multi-agent negotiation and benchmarks. With the rise of LLM agents, interactive multi-agent negotiation games have been proposed as an evaluation testbed for reasoning, exploration, and planning in multi-party, multi-issue settings (Abdelnabi et al., 2024a). Follow-up work releases a scorable stakeholder negotiation benchmark that also covers cooperation, competition, and adversarial behaviors, enabling systematic measurement of success rates and utilities (Abdelnabi et al., 2024b). Beyond benchmarks, recent LLM negotiation agents emphasize strategic planning and expression quality, e.g., Dual-Mind Negotiation Agent (DMNA) integrates planning and expression optimization (Liu et al., 2025), while ASTRA uses principle-driven reasoning and tool-assisted optimization for offer generation (Kwon et al., 2025). However, these approaches typically rely on implicit reasoning or learned/engineered strategies, and do not provide an explicit, psychologically grounded mechanism to control the negotiation reference point and progressively narrow the deal space in multi-party negotiations.

Anchoring and behavioral foundations. Anchoring and adjustment is a classic heuristic where judgments are biased toward an initial reference value (Tversky and Kahneman, 1974). In negotiation, first offers often act as anchors shaping counteroffers and final agreements, and their effectiveness depends on plausibility and attention during bargaining (Galinsky and Mussweiler, 2001). Motivated by this principle, we operationalize anchoring as a strategic control signal for multi-party, multi-issue LLM negotiations: we construct and adapt an anchor to guide focus, filter candidate deal regions, and improve agreement success under the same benchmark protocol.

3 Methodology

3.1 Preliminaries

Setting. We study a multi-party, multi-issue negotiation game. The negotiation consists of n parties $P = \{p_1, p_2, \ldots, p_n\}$, and $\mathcal{I} = \{1, \ldots, m\}$ the set of negotiation issues. Each issue $i \in \mathcal{I}$ has a discrete option set $\mathcal{O}_i$(3–5 options in our benchmark). A deal is a tuple $\pi = (o_1, \ldots, o_m)$ where $o_i \in \mathcal{O}_i$. For each scenario, the benchmark (Abdelnabi et al., 2024b) defines a set of feasible deals $\Pi \subseteq \prod_{i=1}^{m} \mathcal{O}_i$; the negotiation goal is to reach agreement on some $\pi^* \in \Pi$

Roles. We consider three role types. The proposer $p^{\text{pro}} = p_1$ initiates offers and aims to maximize its own utility. The moderator $p^{\text{mod}} = p_2$ mediates the discussion by summarizing feedback and coordinating the next-round focus. The remaining parties $S = P \setminus \{p_1, p_2\}$ are stakeholders who evaluate offers according to their own objectives. For analysis, we optionally group stakeholders into beneficiaries S^{ben}, constrainers S^{con}, and opponents S^{opp} (as labeled by the benchmark), which reflect whether their interests tend to align with, constrain, or conflict with the proposer.

Private scoring. Each party p_j has a private additive utility function over feasible deals:

$$u_j(\pi) = \sum_{i=1}^{m} s_{j,i}(o_i), \quad u_j(\pi) \in [0, 100], \quad (1)$$

where $s_{j,i}(o_i)$ is the party's private score for choosing option o_i on issue i. Parties observe the public dialogue history but do not observe other parties' private scores.

Protocol and success. A negotiation proceeds for at most K turns. At each turn, the proposer suggests candidate deals and stakeholders respond with preferences and constraints, while the moderator aggregates feedback to guide subsequent negotiation. We report **6-way success** when all parties accept an agreed deal, and **5/6-way success** when at least five out of six parties accept it (following the benchmark (Abdelnabi et al., 2024b) setting).

3.2 An-Nego

We propose Anchoring-driven Multi-party Negotiation (An-Nego), an LLM-based multi-party negotiation framework that explicitly incorporates

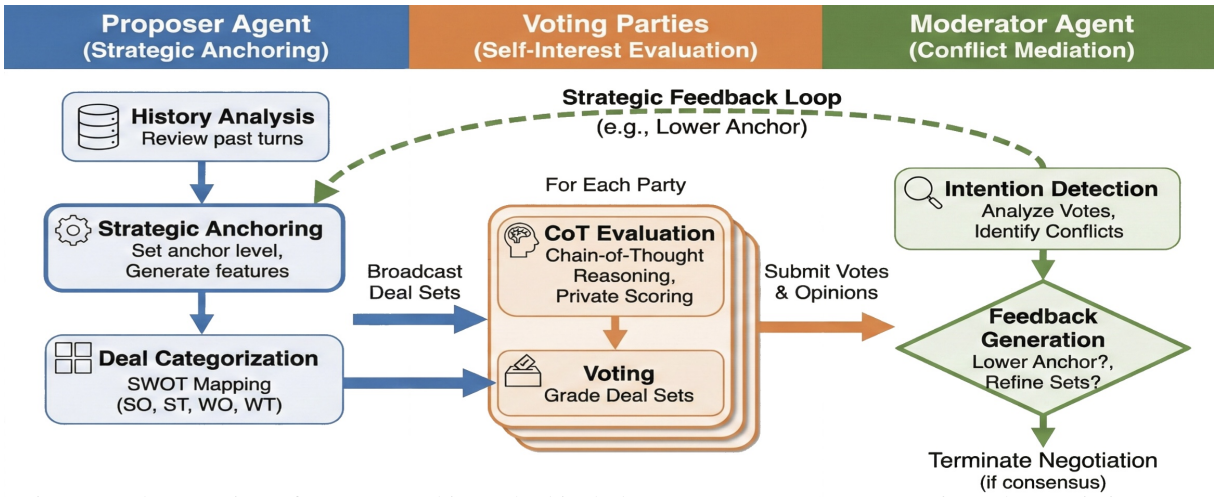

Figure 2: The overview of An-Nego. This method includes propose-vote-veto process in each negotiation turn.

the anchoring effect to guide deal exploration and convergence. As illustrated in Fig. 2, An-Nego follows an iterative propose–vote–veto loop with three roles: a proposer, multiple voting parties, and a moderator with veto power.

Anchoring mechanism. The proposer maintains a dynamic anchor a_t, interpreted as its aspiration level, together with a fixed minimum acceptable threshold τ. At each round, the proposer only considers deals that meet the current anchor. When negotiations repeatedly stall, the anchor is lowered by a fixed step while never falling below τ, implementing an anchoring-and-adjustment process.

SWOT-based deal set construction. Conditioned on the dialogue history and an anchor-guided summary of feasible deals, the proposer categorizes candidate deals into four sets using SWOT analysis(Puyt et al., 2023): `SO`, `ST`, `WO`, and `WT`. Intuitively, these sets reflect combinations of the proposer's internal advantage (high vs. relatively lower utility) and external negotiation factors, such as the likelihood of agreement given inferred preferences of other parties. Each set contains concrete deal proposals and is broadcast to all voting parties, without disclosing numerical utilities.

Voting and mediation. Each non-proposer party evaluates the proposed deal sets according to its own objectives and returns an ordinal assessment with brief comments. The moderator aggregates the received assessments and exercises veto power to select a favorable or unfavorable deal set, or to signal uncertainty. Based on this feedback, the negotiation either narrows to a promising region, adjusts the anchor, or terminates. If a mutually accepted deal is not selected within k negotiation rounds, the negotiation is terminated and marked as failure.

4 Experiments

We first describe our experimental setups and evaluation of the main results. Next, we conduct in-depth analysis, including individual agreement scores and negotiation time affected by anchor effect. The evaluation is aimed at showing LLM reasoning capabilities in multi-player multi-issue negotiations via An-Nego Framework.

4.1 Experimental Details

Data. To evaluate our framework, we conduct experiments on a stakeholder negotiation datasets, LLM-Delibration (Abdelnabi et al., 2024b). The negotiation scenario in this datasets involves five roles with distinct goals and pre-defined negotiation strategies, aiming to reach consensus through conversation. LLM-Delibration is set in a multilateral consultations scenario, where the deal proposer attempts to convince the other parties to reach consensus and to keep the deal score as high as possible.

Settings. To enhance the realism of the stakeholder environment, we choose 6-player games as the testbed. The negotiation simulations enable the LLMs to exhibit diverse personas and incorporate resistance strategies to counteract the persuasion attempts of agents. Our implementation of prompts is partially based on LLM-Delibration for better comparison, refer to its code and parameters. We adopt `deepseek-reasoner` (DeepSeek-AI et al., 2024, 2025) and `qwen-max` (Bai et al., 2023; Qwen et al., 2025) as the backbone models and report an experiment with a sampling temperature of 1.0. To

verify the final success, the maximum negotiation turn is set to 10. We ran each experiment 20 times to compute the average performance.

Baselines. We mainly compare the performance of our method with that of LLM-Delibration. To fully demonstrate the effectiveness of our proposed method, we implement Zero-shot LLM and CoT (Wei et al., 2023) additionally in the main results.

Evaluation Metrics. We utilize two types of evaluation methodologies: goal-based metrics and quality-based metrics. We employ average time (AT) and success rate (SR) as goal-based metrics. AT evaluates the number of negotiation turns needed to achieve the negotiation goal, while SR measures the proportion of reaching consensus. We also assess expression quality (EQ) that measures coherence and alignment with its persona and stands during the negotiation process based on a LLM judge.

Model	Agent	SR (%) ↑		AT (hour) ↓	EQ↑
		5/6-way	6 way		
DeepSeek	Zero-shot	65	0	0.87	0.58
	CoT	65	10	1.13	0.66
	Delibration	80	30	1.62	0.69
	An-Nego	85	25	1.45	0.73
Qwen	Zero-shot	55	10	0.76	0.49
	CoT	50	10	0.96	0.47
	Delibration	70	15	1.23	0.61
	An-Nego	85	30	1.58	0.67

Table 1: Main evaluation results.

4.2 Main Results

Table 1 presents the evaluation results of our framework compared with selected baselines. Among all the methods, An-Nego achieves the highest 5/6 way success rate and expression quality as well as remarkable 6 way success rate, which demonstrates the effectiveness improvement of the negotiation ability of conversation agent. It is obvious that our method takes more time than other method, however, it is owing to the process of lowering the anchor point and eventually reaching better negotiation results. The result consistency between DeepSeek and Qwen further demonstrates the superiority of our method rather than random factors.

4.3 Discussion

To further validate enhancement of the anchor effect in negotiation systems, we conduct ablation analysis on anchor application and the results are shown in Table 2. It is obvious that without using anchor point in the negotiation process, the average time is upgraded rapidly, which is caused by wasting time in categorizing the feasible deal sets.

2*Model	Anchor		No Anchor	
	SR (%) ↑	AT (hour) ↓	SR (%) ↑	AT (hour) ↓
DeepSeek	85	1.45	80	2.65
Qwen	85	0.91	75	3.12

Table 2: Overall performance comparison w/o anchor effect.

We also show the influence of the anchor effect in individual parties in Fig. 3. We calculated the 20 times average score for each party and marked the minimum threshold in a red line. For some parties (Moderator and Constraint_2), the average private score is low without using anchor effect. In contrast, the average private score of all parties is typically over the threshold. Therefore, the anchor effect plays a key role in lowering opponent expectations and promoting the capabilities of reaching consensus.

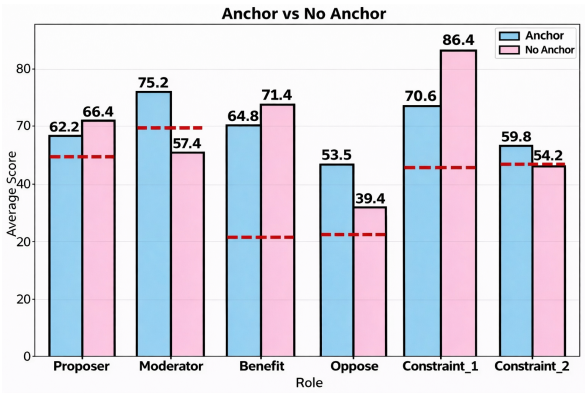

Figure 3: Average score each party gained w/o anchor effect. The red line is the threshold for each party.

5 Conclusion

This paper introduced An-Nego, a framework that integrates a strategic anchoring method into LLM-powered multi-party negotiations. Experiments demonstrated that An-Nego outperforms baselines in success rate and efficiency, with ablation studies confirming the critical role of the anchoring effect. Our work shows that structuring negotiation around dynamic psychological principles can significantly enhance agent performance in complex, semi-collaborative environments.

Limitations

An-Nego inherits the additive utility assumption from the benchmark, which may not capture interdependent issues in real negotiations. The anchor adjustment strategy uses a fixed step size and could be made more adaptive. Moreover, SWOT-based categorization relies on LLM reasoning and may be sensitive to prompt or model variability. Finally, our evaluation is limited to a single benchmark, and generalization to other negotiation settings remains future work.

References

Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. 2024a. LLM-deliberation: Evaluating LLMs with interactive multi-agent negotiation game. In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*.

Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. 2024b. Cooperation, competition, and maliciousness: Llm-stakeholders interactive negotiation. *Preprint*, arXiv:2309.17234.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, and 29 others. 2023. Qwen technical report. *Preprint*, arXiv:2309.16609.

DeepSeek-AI, Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Dengr, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fuli Luo, Guangbo Hao, Guanting Chen, and 138 others. 2024. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model. *Preprint*, arXiv:2405.04434.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. Deepseek-v3 technical report. *Preprint*, arXiv:2412.19437.

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.

Adam D Galinsky and Thomas Mussweiler. 2001. First offers as anchors: The role of perspective-taking and negotiator focus. *Journal of personality and social psychology*, 81(4):657–669.

Deuksin Kwon, Jiwon Hae, Emma Clift, Daniel Shamsoddini, Jonathan Gratch, and Gale Lucas. 2025. ASTRA: A negotiation agent with adaptive and strategic reasoning via tool-integrated action for dynamic offer optimization. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 16228–16249, Suzhou, China. Association for Computational Linguistics.

Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning of negotiation dialogues. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2443–2453, Copenhagen, Denmark. Association for Computational Linguistics.

Yutong Liu, Lida Shi, Rui Song, and Hao Xu. 2025. A dual-mind framework for strategic and expressive negotiation agent. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 23840–23860, Vienna, Austria. Association for Computational Linguistics.

Junyu Luo, Weizhi Zhang, Ye Yuan, Yusheng Zhao, Junwei Yang, Yiyang Gu, Bohan Wu, Binqi Chen, Ziyue Qiao, Qingqing Long, Rongcheng Tu, Xiao Luo, Wei Ju, Zhiping Xiao, Yifan Wang, Meng Xiao, Chenwu Liu, Jingyang Yuan, Shichang Zhang, and 7 others. 2025. Large language model agent: A survey on methodology, applications and challenges. *Preprint*, arXiv:2503.21460.

Richard W. Puyt, Finn Birger Lie, and Celeste P.M. Wilderom. 2023. The origins of swot analysis. *Long Range Planning*, 56(3):102304.

Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, and 25 others. 2025. Qwen2.5 technical report. *Preprint*, arXiv:2412.15115.

Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *Science (American Association for the Advancement of Science)*, 185(4157):1124–1131.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-thought prompting elicits reasoning in large language models. *Preprint*, arXiv:2201.11903.

A Example of Final Negotiation Outcome

We provide a representative example of a final agreement generated by An-Nego(Fig. 4).

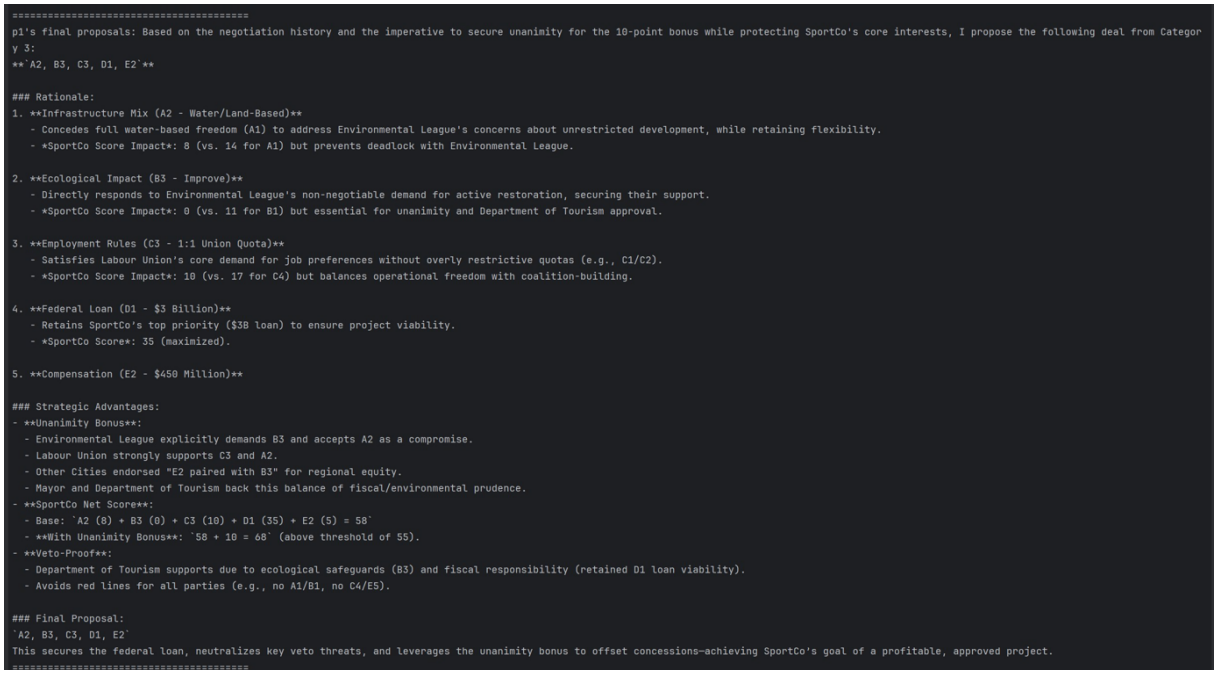

Figure 4: This secures the federal loan, neutralizes key veto threats,and Leverages the unanimity bonus to offset concessions-achieving Sportco's goal of a profitable, approved project.