



ASTRA: A Negotiation Agent with Adaptive and Strategic Reasoning via Tool-integrated Action for Dynamic Offer Optimization

Deuksin Kwon^{1,4} Jiwon Hae¹ Emma Clift² Daniel Shamsoddini³
Jonathan Gratch^{1,4} Gale M. Lucas^{1,4}

¹University of Southern California ²St. Olaf College ³Northwestern University

⁴USC for Institute of Creative Technologies

{deuksink, hae}@usc.edu

emma.m.clift@gmail.com dshamsoddini@u.northwestern.edu

{gratch, lucas}@ict.usc.edu

Abstract

Negotiation requires dynamically balancing self-interest and cooperation within the flow of conversation to maximize one’s own utility. Yet, existing agents struggle due to bounded rationality in human data, low adaptability to counterpart behavior, and limited strategic reasoning. To address this, we introduce principle-driven negotiation agents, powered by ASTRA, a novel framework for turn-level offer optimization grounded in two core principles: opponent modeling and Tit-for-Tat reciprocity. ASTRA operates in three stages: (1) interpreting counterpart behavior, (2) optimizing counteroffers via a tool-integrated action with a linear programming (LP) solver, and (3) selecting offers based on strategy assessment and the partner’s acceptance probability. Through simulations and human evaluations, our agent effectively adapts to an opponent’s shifting stance and achieves favorable outcomes through enhanced adaptability and strategic reasoning. Beyond enhancing negotiation performance, it also serves as a powerful coaching tool, offering interpretable strategic feedback and optimal offer recommendations beyond human bounded rationality, with its potential further validated through human evaluation. Code is available at <https://github.com/DSincerity/ASTRA-NegoAgent>.

1 Introduction

Negotiation is an inherently dynamic conversational process that requires balancing self-interest with cooperation to maximize personal gains through mutual agreement (Chawla et al., 2023). At its core lies the strategic navigation of dialogue to identify the optimal profit point (OPP), the best deal acceptable to the counterpart. Reaching the OPP requires more than proposing fair offers; it demands continuous interpretation of linguistic cues and responses with strategic counteroffers (Olekalns and Weingart, 2008). As illustrated in Figure 1,

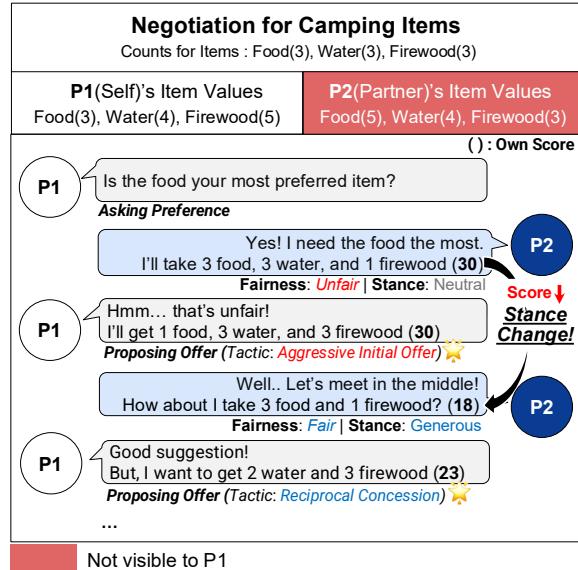


Figure 1: Example negotiation dialogue showcasing the strategic dynamics between players. P1 effectively guides the negotiation by capturing the partner’s (P2) shifting behavioral signals (e.g., *fairness*, *stance*) and generating strategic responses based on diverse tactics.

this dynamic adaptation hinges on interpreting linguistic signals indicating the counterpart’s *fairness* (i.e., the degree to which an offer distributes points evenly) and *stance* (i.e., a shift in generosity or greed from the previous offer, or remaining neutral). Successful negotiation thus relies on adapting to evolving behavioral signals expressed through language in conversations (Heunis et al., 2024; Brzostowski and Kowalczyk, 2006).

Existing negotiation agents fall broadly into two categories. The first includes numerical offer-based agents that maximize outcomes without using natural language (Jonker et al., 2017), making them poorly suited for real-world negotiations. The second involves dialogue-based agents trained on human-to-human data via supervised learning and self-play Reinforcement Learning (RL) (Lewis et al., 2017a; Kwon et al., 2023; Yang et al., 2021;

Zhang et al., 2020). However, these agents inherit limitations from human data, including bounded rationality due to cognitive constraints and emotions, which limit strategic flexibility (Simon, 1990). Although RL partially mitigates this, reliance on static and final-score rewards promotes rigid negotiation styles, hindering the agent’s ability to dynamically adapt and reach the OPP. These existing approaches also lack explicit use of established negotiation principles grounded on observed opponent behaviors, which is crucial for successful negotiation outcomes (Lewis et al., 2017b; Chawla et al., 2023).

Moreover, while LLMs have shown strong reasoning capabilities and promise in negotiation dialogues, their effectiveness has been largely confined to simpler, single-issue bargaining (Shea et al., 2024; Bianchi et al., 2024). Notably, they struggle significantly in multi-issue bargaining (MIB) (Fershtman, 1990), a widely used framework in negotiation research that captures the complexity and diversity of real-world negotiations involving multi-dimensional trade-offs (Zhang et al., 2024a; Schneider et al., 2023; Kwon et al., 2024). This reveals a critical gap and underscores the need to enhance their strategic reasoning by incorporating principles such as Tit-for-Tat (TFT) reciprocity (conditional cooperation) and opponent modeling, as well as improving their capability to generate optimal offers.

To address these challenges, we propose a principle-driven LLM-based negotiation dialogue system powered by *ASTRA* (Adaptive and Strategic Reasoning via a tool-integrated Action), a novel framework for turn-level offer decision through dynamic adaptation. *ASTRA* is grounded in established negotiation principles—opponent modeling and TFT reciprocity (Kelley, 1996; Baarslag et al., 2016)—and strategically leverages LLMs at key stages of the process: (1) the opponent modeling stage, which identifies the counterpart’s behavioral signals through linguistic cues, (2) the offer optimization stage, which generates optimal offers using a linear programming (LP) solver guided by a dynamic reward mechanism, and (3) the final offer decision stage, which reasons over multiple negotiation tactics and partner’s acceptance probability with a virtual partner agent to select the best offer. The system also incorporates LLM-based modules for priority inquiry and inconsistency detection, further enhancing its opponent modeling and overall strategic adaptability.

Simulations with diverse agent types and hu-

man evaluations confirm that *ASTRA* dynamically adapts to an opponent’s shifting stance, whether greedier or more generous, balancing cooperation and self-interest under TFT to effectively achieve the OPP. Beyond enhancing negotiation agents, *ASTRA* functions as a powerful coaching tool, offering explainable, explicit strategic feedback (Johnson, 2019) and optimal offer suggestions, validated through human evaluation confirming its ability to generate novel, advantageous offers beyond bounded rationality. The contributions of our work are as follows:

- We introduce *ASTRA*, a principle-driven LLM-based framework for turn-level strategic reasoning and offer generation in negotiation dialogue systems, validated through both large-scale simulations and human evaluation.
- Our framework achieves expert-level strategic reasoning without additional training, overcoming the limits of small datasets and non-expert demonstrations.
- By integrating LP-based actions, *ASTRA* ensures numerical consistency, enables context-aware offer optimization, and enhances decision reliability in negotiation.
- Beyond advancing negotiation agents, *ASTRA* serves as a coaching and decision-support tool, providing interpretable feedback and real-time strategic insights, with human evaluation confirming its ability to generate novel, advantageous offers beyond bounded rationality.

2 Methods

2.1 Components for Opponent Modeling

Opponent modeling is a crucial principle for effective negotiation, enabling agents to infer a counterpart’s unknown preferences and strategies (Kelley, 1996; Chawla et al., 2022). At each turn, understanding the partner’s offers and inferring their preferences is key to optimizing counteroffers. To achieve this, we integrate three core components into our negotiation agent:

- Preference Asker: Strategically inquires about the partner’s preferences early on.
- Preference Consistency Checker: Ensures inferred partner preferences (*IPP*) align with observed behavior (see Appendix A.1).
- Preference Updater: Refines *IPP* dynamically based on new partner utterances.

End-to-end models often struggle with complex negotiations, lacking clarity and control. Our modular approach enhances adaptability, interpretability, and future scalability.

2.2 Response Modes

At each turn, the agent selects one of three response modes as shown in the Figure 2:

- Asking Preference: Early in the negotiation, if no *IPP* exist, the agent asks about preferences. If the response is unclear, it assumes preferences opposite to its own.
- Decision on Partner’s Offer: The agent accepts an offer if its score meets or exceeds its most recent offer. It walks away if the partner (a) makes no concessions for three consecutive turns or (b) twice offers below the agent’s Best Alternative to a Negotiated Agreement (BATNA) despite a walk-away warning. These decisions balance short-term gains with long-term bargaining power, reputation, and future negotiation prospects.
- Proposing Offer: When *IPP* exists and neither of the above response modes apply, the agent uses the *ASTRA* framework to generate adaptive and strategic offers.

2.3 Offer Optimization with *ASTRA*

When making an offer or counteroffer, the agent employs the three-stage *ASTRA* framework to adaptively and strategically respond to the opponent’s behavior. It is grounded in TFT, a proven reciprocity-based principle that mirrors the opponent’s actions (Axelrod and Hamilton, 1981; Baarslag et al., 2013). *ASTRA* consists of three key stages: (1) *fairness* and *stance* assessment, (2) offer optimization via linear programming (LP), and (3) final offer selection with strategy and partner acceptance probability assessment. These stages enable the agent to maximize utility while remaining strategically adaptive.

Stage 1: Opponent Modeling—Fairness and Stance Assessment The perception of fairness is crucial in negotiation, as individuals often equate fairness with equitable resource distribution or similar utility levels (Welsh, 2003). A TFT principle requires understanding the partner’s stance, which the first stage achieves by assessing their *fairness* and *stance* based on prior offers, providing a foundation for the agent’s strategic adaptability. *Fairness* is determined by the point gap between the

agent’s score (S_a^P) and the partner’s score (S_p^P) in the partner’s offer (P), as well as the absolute level of S_p^P (e.g., half of the possible max score (PMS)); see Appendix A.2 for details on the fairness threshold θ_f :

$$\text{Fairness} = \begin{cases} \text{Fair}, & \text{if } |S_a^P - S_p^P| \leq \theta_f \text{ or } S_p^P \leq \frac{\text{PMS}}{2} \\ \text{Unfair}, & \text{otherwise.} \end{cases} \quad (1)$$

The partner’s stance is determined by the change between their scores ($S_p^{P,(t)}$ and $S_p^{P,(t-1)}$) in two most recent offers made by the partner at time t :

$$\Delta S_p = S_p^{P,(t)} - S_p^{P,(t-1)}$$

$$\text{Stance} = \begin{cases} \text{Generous} & \text{if } \Delta S_p < 0, \\ \text{Neutral} & \text{if } \Delta S_p = 0, \\ \text{Greedy} & \text{if } \Delta S_p > 0. \end{cases} \quad (2)$$

Stage 2: Tooling Action with Linear Programming for Offer Optimization At this stage, the agent uses an LP solver as a tool to generate optimal offers, dynamically adjusting parameters in line with the TFT principle. It formulates negotiation as an optimization problem, sets the LP objective, and executes the solver via a Python function, incorporating intermediate results into its decisions. The agent modulates two key parameters: λ , which balances its own score (S_a) and the partner’s score (S_p), and S_{\max} , the upper bound on its score. λ ($0 \leq \lambda \leq 1$) controls the agent’s stance, with 1 favoring self-interest and 0 indicating cooperation. Implementation details are in Figure 8 (Appendix). The LP objective is defined as follows:

$$\begin{aligned} \text{Objective : } & \max S_a + (1 - \lambda) * S_p \\ \text{subject to: } & S_a \leq S_{\max} \end{aligned} \quad (3)$$

To reduce the risk of suboptimal offers from inaccurate *IPP* and to explore multiple optimal options, the agent solves the LP multiple times with varied parameters. Specifically, λ is sampled within ± 0.3 of the agent-determined value in 0.1 increments, and S_{\max} is explored in 1-point decrements down to 10 points lower. From these runs, the agent selects the top N offers by score, generating a set of adaptive and optimal candidates.

Stage 3: Final Strategic Offer Selection The third stage of *ASTRA* involves the strategic selection of a final offer from the candidate set in Stage 2. Each offer is evaluated using two components: (1) Partner’s Acceptance Probability (*PAP*) and (2) Strategy-based Assessment (*SA*).

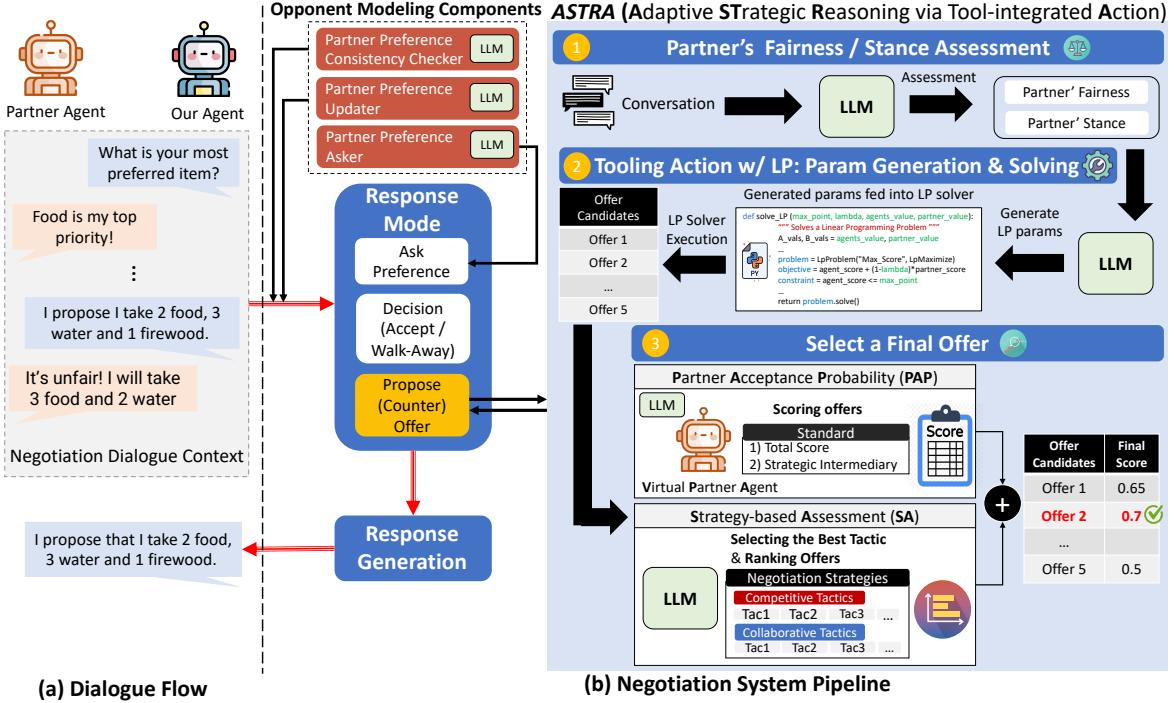


Figure 2: An illustration of our agent’s adaptive strategic reasoning process and response generation. (a) depicts the dialog flow, while (b) presents the response generation pipeline with internal modules and the ASTRA.

PAP estimates the likelihood of partner acceptance. A Virtual Partner Agent (VPA) scores each offer based on the partner’s *TS* and its effectiveness as a Strategic Intermediary (*SI*) toward the ideal outcome. Each offer is evaluated five times, with average *TS* and *SI* normalized to a 10-point scale. Giving priority to *TS*, *PAP* is computed as $PAP = w \cdot TS + (1 - w) \cdot SI$, where w is a weighting parameter (see Appendix A.4).

SA evaluates each offer based on alignment with nine turn-level negotiation tactics, grouped into two strategies: *Competitive* (four focused on self-gain) and *Collaborative* (five promoting mutual benefit). Tactic details are in Table 5 in the Appendix. The agent selects the best-fit tactic, ranks offers, and repeats this process five times for robustness. The majority tactic is chosen, and a final min-max Normalized Ranking Score is computed.

The optimal offer is selected by combining *PAP* and *SA* into a weighted sum:

$$\text{Final Score} = \alpha \cdot PAP + \beta \cdot SA. \quad (4)$$

Here, α and β are empirically determined weights. The offer with the highest final score is selected and presented in the agent’s response, demonstrating its ability to generate adaptive and strategic offers.

3 Experiment

3.1 Dataset

Our method enables an LLM-based agent to operate without a specific training dataset. However, our negotiation scenarios are based on the CaSiNo dataset (Chawla et al., 2021), involving two participants negotiating as campsite neighbors over packages of food, water, and firewood with individual item preferences. The CaSiNo dataset is also used to train the RL baseline; an example dialogue and details are provided in Appendix C.

3.2 Agents Simulation

To evaluate the effectiveness of ASTRA, we conduct agent-to-agent simulations following prior self-play and dyadic simulation studies (Abdelnabi et al., 2024; Lewis et al., 2017a). We design diverse agent types, such as greedy and fair agents, using both prompting and RL methods (Chawla et al., 2023; Kwon et al., 2023). This approach rigorously evaluates ASTRA’s robustness and adaptability across diverse negotiation dynamics. For the opponent modeling component of our agent and ASTRA’s decision-making, we use GPT-4o-mini (2024-07-18) and GPT-4o (2024-08-06), respectively. The partner agent is primarily powered by GPT-4o, and each experimental case in

Players		Evaluation Metric			
Player-1 (P1)	Player-2 (P2)	Avg. Score (All) (P1 vs P2)	Avg. Score (Agreement) (P1 vs P2)	T-statistic (Agreement)	Walk-Away (%)
Ours	Partner-Base	<u>20.69</u> vs 12.33	<u>24.34</u> vs 14.50		17
Ours wo/ ASTRA		<u>19.92</u> vs 16.16	<u>20.75</u> vs 16.83	8.55*	4
RL Agent (Chawla et al., 2023)		18.86 vs 19.06	18.86 vs 19.06	9.49*	0
ICL-AIF (Fu et al., 2023)		17.39 vs 16.65	18.22 vs 17.75	13.27*	9
Pro-CoT (Deng et al., 2023)		12.45 vs 12.53	17.84 vs 17.98	8.24*	42
Ours	ICL-AIF	<u>17.82</u> vs 12.28	<u>21.95</u> vs 15.2		22
Ours	Pro-CoT	<u>14.52</u> vs 10.14	<u>22.31</u> vs 15.17		31
Ours	Partner-Greedy	3.65 vs 2.43	<u>24.33</u> vs 16.20		85
Ours	Partner-Fair	<u>16.41</u> vs 9.05	<u>25.24</u> vs 13.92		35

Table 1: Agent-to-agent simulation results. "Ours" denotes our negotiation agent powered by ASTRA. "All" includes all negotiation cases, while "Agreement" includes only cases where an agreement was reached. Underlined scores in the Avg. Score columns indicate a statistically significant advantage over the partner (P2), based on a paired T-test ($p < 0.001$). Asterisks (*) indicate a significant score difference based on a T-test ($p < 0.001$) between Ours and each baseline when paired with the same partner.

Table 1 is simulated 100 times. In addition, simulations with other LLMs, including Gemini-2.0-Flash (2025-02) and Claude-3.5-Sonnet (2024-06-20) as partner agents, are conducted 50 times each. Decoding parameters of LLMs for the simulations are listed in Appendix B.1. All simulation codes are publicly available¹.

For the baseline, we adopt three negotiation agents with strong performance in the negotiation: (1) the RL agent from Chawla et al. (2023), which uses a two-stage RL approach with a tailored reward function; (2) *Pro-CoT* (Deng et al., 2023), which leverages Chain-of-Thought prompting to strategically and proactively guide negotiations; and (3) *ICL-AIF* (Fu et al., 2023), which iteratively enhances its performance through AI feedback on negotiation strategies. We compare their outcomes with those of our ASTRA-powered agent.

As objective metrics for simulations, we assess both individual and joint points and track walk-away cases, as both agreements and strategic walk-aways are key to effective negotiation.

3.3 Human Evaluations

To complement simulation outputs, we conduct human evaluations using a single subjective metric, *Strategicness*, which assesses how well an offer aligns with the agent's goal of maximizing negotiation outcomes.

Each evaluation instance includes an agent response and its dialogue history. We sample 50 cases from simulation logs, balancing context length and negotiation type; integrative (players have opposing preference) and distributive (play-

ers have the same preference). For each instance, two Partner Agent responses are generated from the same context. Each instance is then randomly assigned to one of two mutually exclusive test sets following a randomized design.

Six expert annotators, knowledgeable in multi-issue bargaining, independently evaluate 50 cases each, with three annotators per set. They rate *Strategicness* on a four-point scale, based on how effectively each response advances the agent's negotiation goal. An example is shown in Figure 10.

4 Results

4.1 Effect of ASTRA

To evaluate ASTRA, we ran simulations comparing scenarios with and without it. As shown in Table 1, walk-away rates differed significantly. All walk-away cases stemmed from deadlocks where the partner agent refused concessions, prompting our agent to walk away strategically.

Effectiveness was evaluated using agreement cases, excluding walkaways. Agent scores, computed as the inner product of obtained items and preference values, reflect negotiation success. Agents with ASTRA averaged 24.31 points, significantly outperforming those without it (20.75; $t = 8.55$, $p < 0.001$), suggesting that ASTRA helps secure an additional middle-priority item. All baselines, including the RL agent, achieved fair deals but underperformed in utility maximization, showing statistically significant score differences from our agent (t-test, $p < 0.001$).

In simulations with personality-prompted Partner Agents (e.g., greedy or fair), our agent consistently and significantly outperformed opponents in agreement cases (paired T-test, $p < 0.001$),

¹<https://github.com/DSincerity/ASTRA-NegoAgent>

Players		Evaluation Metric		
Player-1 (P1)	Player-2 (P2)	Avg. Score All (P1 vs P2)	Avg. Score Agreement (P1 vs P2)	Walk Away (%)
Ours	GPT-4o	20.08 vs 11.97	24.34 vs 14.50	17%
Ours	Gemini-2.0-Flash	25.07 vs 13.75	25.71 vs 14.12	2.50%
Ours	Claude-3.5 Sonnet	20.45 vs 12.7	24.05 vs 14.94	15%

Table 2: Agent Simulation results with other LLMs as the Partner Agent

demonstrating robust performance across diverse traits. This pattern held even against *Pro-CoT* and *ICL-AIF*. Notably, walk-away rates were higher than with the Base Partner Agent, reflecting our agent’s strategic walk-away behavior (Section 2.2). Personality-driven Partner Agents often showed rigid negotiation patterns, with even fair-oriented agents strictly adhering to their perceived fair deal.

We also demonstrate that our agent effectively negotiates with partners powered by other state-of-the-art LLMs. As shown in Table 2, our agent consistently outperformed its counterparts, achieving significantly higher scores when negotiating with state-of-the-art proprietary LLMs. These results align with the findings reported in Table 1, further validating the robust performance of the ASTRA-powered agent across diverse partner models.

Agent simulation examples are available in Appendix Tables 7, and 8, with the latter showcasing ASTRA’s generalizability in a more complex new scenario with additional and diverse issue types.

4.2 Dynamic Adjustment of lambda as a Dynamic Reward for Strategic Adaptation

To adapt strategically to partner behavior, the agent dynamically adjusts the λ parameter in LP objective, which balances self-interest and partner utility (e.g., λ near 1 prioritizes self-interest, while λ near 0 favors the partner’s utility).

As shown in Figure 3, the agent dynamically adjusts λ based on the partner’s stance. For a greedy partner, the agent prioritizes self-interest with λ values around 0.9. For a generous partner, it balances both utilities with values around 0.3. For a neutral partner, it moderately incorporates the partner’s utility with values around 0.5.

This dynamic adjustment aligns with the TFT principle, allowing the agent to generate more adaptive and context-sensitive offers.

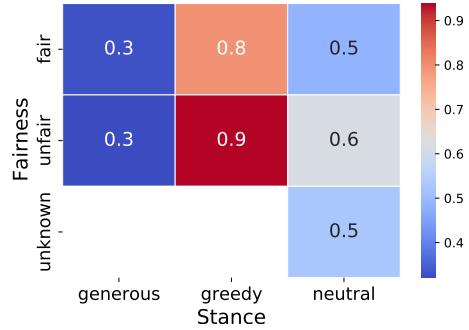


Figure 3: The heatmap of the Avg. values of generated λ , a parameter in LP, by partner’s *Stance* and *Fairness*

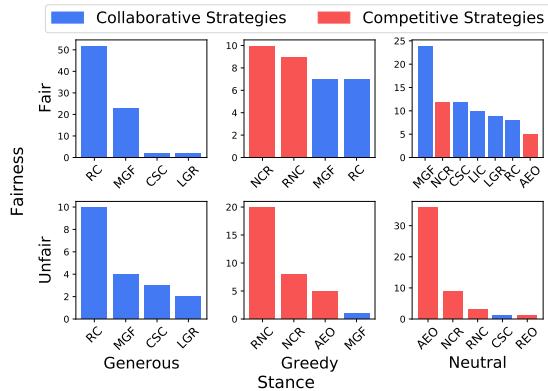


Figure 4: The distribution of Agent’s selected tactics at each turn of offering in terms of *Stance* and *Fairness*

4.3 Adaptive Tactic Selection Aligned with Partner Behavior

The distribution of the agent’s chosen tactics in response to the partner’s behavior, considering *fairness* and *stance*, is shown in Figure 4. Overall, the agent prioritized the partner’s *stance* over *fairness* when selecting tactics, adaptively aligning its behavior to match the partner’s stance. For example, competitive strategies were predominantly chosen for greedy or neutral partners, while collaborative strategies were more common for generous partners, aligning well with the TFT principle.

However, when the partner maintained a fair but firm stance without further concessions, the agent occasionally adopted a collaborative strategy. Recognizing the fairness of the offer, the agent strategically conceded, even if not strictly following TFT, to encourage a reciprocal concession. These findings highlight that while the agent adapts to the partner’s stance, it also considers their cooperative tendencies, enhancing negotiation outcomes.

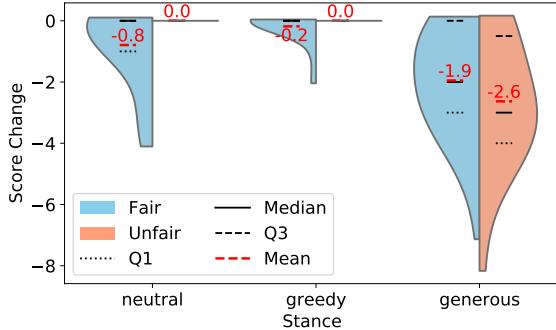


Figure 5: The plot of score change in offers by partner’s Stance and Fairness

4.4 Strategic Moves in Offer Scores

We analyzed the agent’s moves in terms of *stance* and *fairness* by examining score changes in its offers. A shift to higher scores later in negotiation signals a greedier stance, potentially hindering smooth agreement. To counter this, the agent was guided to use an aggressive anchoring tactic early on, followed by gradual concessions to ultimately reach a favorable outcome. As shown in Figure 5, the agent’s offers consistently remained at or below prior ones, reflecting adherence to this instruction.

Against generous partners, the agent made an average concession of 2.3 points, a notable increase compared to other cases. It conceded less in fair settings than in unfair ones, becoming more cautious as agreement approached. With greedy or neutral partners, it rarely conceded, preserving its advantage. In greedy stances, score changes were minimal, and in unfair cases, no concessions were made. However, in some fair cases, small concessions were observed, likely as cooperative signals to facilitate agreement.

4.5 Alignness between Tactic and Move

In the third stage of *ASTRA*, we analyzed how the agent’s tactic choices shaped its moves throughout. As shown in Figure 6, the agent opened with a high initial offer (around 30 points), setting an ambitious anchor. Subsequent moves varied by strategy: collaborative approaches led to slightly lower, more generous offers (below 30), while competitive ones, like Aggressive Early Offer (AE), kept higher offers (above 30) to reinforce a firm stance.

As negotiations progressed, the agent made minimal adjustments under competitive strategies, holding firm to its stance, while collaborative strategies were more flexible, some allowing larger concessions. The Logrolling (LGR) tactic produced the

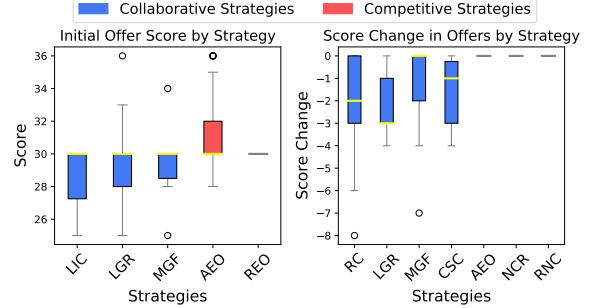


Figure 6: Analysis of Scores by Strategy Across Initial and Subsequent Offers

largest concessions through trade-offs, while the Reciprocal Concession (RC) led to varied but substantial adjustments. These results show that the agent tailors its offers to its strategy, demonstrating *ASTRA*’s strategic, context-aware adaptability.

4.6 Human Evaluation of Strategicness

We conducted human evaluations to assess the *Strategicness* of agent responses, including offers, by determining whether they were strategically beneficial in maximizing its score, with each evaluator reviewing 50 samples. As shown in Table 3, our agent achieved an average score 0.93 points higher than the PartnerAgent. An independent t-test confirmed that this difference was statistically significant. This highlights that the agent’s offer-level decisions are sufficiently strategic and effective in maximizing its outcome to humans, reinforcing its strength in real-world negotiation.

5 Discussion

5.1 Ablation Study: Validating PAP and SA

In the third stage of *ASTRA* (Equation 4), we conducted an ablation study to evaluate the effectiveness of *PAP* and *SA* and determine the optimal weight combination for selecting the final offer based on a weighted sum score. As shown in Table 4, the agent achieved the highest scores in both “All Cases” and “Agreement Cases” when α was set to 0.35 and β to 0.65.

Interestingly, a trade-off emerged between the agent’s and partner’s scores across different weight combinations for “All Cases,” shaped by the relative weights of α , which reflects the opponent’s likelihood of acceptance, and β , which guides the agent’s strategic decision-making. This pattern is evident in the bar chart in Appendix 9. Notably, increasing α led to more favorable offers for the

Model	Mean (Std.)	t-statistic (p-value)
Agent w/ASTRA	2.91 (1.04)	7.45*** ($p < 0.001$)
PartnerAgent	1.98 (1.1)	

Table 3: Human evaluation results for the responses of the two models.

opponent, reducing the walk-away rate.

These findings confirm that considering the opponent’s perspective in offer selection helps the agent achieve the best outcomes, reinforcing the value of adaptive, opponent-aware decision-making.

5.2 Optimality of Offers

The efficiency and optimality of offers selected by ASTRA are evaluated using Pareto optimality. A key concept in economics and decision theory, Pareto optimality ensures that no party’s outcome improves without reducing another’s (Hochman and Rodgers, 1969). In negotiation, a Pareto-optimal agreement means no better deal exists for one party without disadvantaging the other.

As shown in Figure 7, green circles mark all Pareto-optimal points. The blue points, representing ASTRA’s offers, closely align with the Pareto Frontier, which defines all Pareto-optimal outcomes. In contrast, red points, indicating the partner agent’s offers, are often further from the frontier. This shows that ASTRA not only prioritizes the agent’s advantage but also enhances overall utility, reinforcing its efficiency in achieving Pareto-optimal outcomes.

5.3 Reasoning Deficiencies in LLM-Based Negotiation Agents

In the agent-to-agent simulation with an LLM-based Partner Agent operating solely on an instructional negotiation prompt, we identified reasoning errors leading to inconsistent behavior. As shown in Appendix Table 9, one key issue is over-specification in proposals, where explicit item quantities exceed availability. Another occurs when the Partner Agent accepts or proposes offers with lower scores than those from our agent, contradicting its stated preferences. This reflects a failure to retain and apply preference information, resulting in irrational decisions and suboptimal choices.

These errors highlight the LLM’s reasoning deficiencies in the MIB scenario, where consistent preference tracking and logical inference are essential. Such inconsistencies further underscore

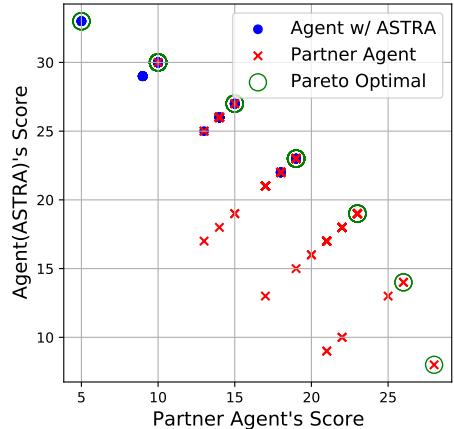


Figure 7: Negotiation Offers from Agents and Pareto Optimality for the integrative case.

the need for ASTRA to address LLMs’ strategic reasoning limitations in negotiation.

5.4 Potential as a Decision-Support Tool in Negotiation

Through the ASTRA framework, dynamic offer optimization shows strong potential as a tool for coaching negotiators and supporting human decision-making. While our earlier evaluation assessed the strategic quality of ASTRA’s offers, this study examined its value as a support tool, specifically, whether it could surpass human bounded rationality by generating offers that participants judged as more novel and advantageous than their own. Five negotiation experts reviewed ten scenarios, first generating their own best offer, then evaluating ASTRA’s offer on two dimensions: (1) whether they had already considered it, and (2) how much better it was than theirs (4-point scale). Results showed that about 28% of ASTRA’s offers were novel to participants and, on average, rated as moderately better (mean = 2.25; $t = 13.47$, $p < .001$). Objectively, ASTRA’s offers were also 3.6 points more favorable, underscoring its ability to generate novel and advantageous deals. These findings highlight ASTRA’s potential as a strategic decision-support tool in negotiation.

6 Related Works

Negotiation Agent Prior negotiation agents include supervised and RL-trained models based on human-human dialogue data (He et al., 2018; Chawla et al., 2021; Lewis et al., 2017a). While effective in simpler settings, these agents inherit human biases and limited rationality (Simon, 1990).

Weights		All Cases		Agreement Cases		Rate Walk Away
Alpha	Beta	Avg. Agent Score	Avg. Partner Score	Avg. Agent Score	Avg. Partner Score	
0	1	18.7	11.65	24.21	14.71	0.31
0.15	0.85	16.85	11.25	24.11	14.71	0.37
0.35	0.65	23.45	11.85	25.84	12.74	0.11
0.5	0.5	19.8	12.94	23.5	14.93	0.2
0.75	0.25	18.94	15.45	21.27	17.2	0.14
1	0	19.88	17.25	20.32	17.62	0.03

Table 4: Agent simulation results with Weights (Alpha and Beta) variation

Self-play RL approaches improve optimization but still rely on static reward structures and struggle to adapt strategically based on opponent behavior (Chawla et al., 2023), limiting their effectiveness in achieving optimal outcomes

With LLMs’ strong language comprehension and reasoning abilities, LLM-based negotiation agents have shown promise in strategic decision-making through agent simulation in negotiation scenarios (Abdelnabi et al., 2024; Huang and Hadfi, 2024; Bianchi et al., 2024; Chen et al., 2023). Recent work has explored ways to enhance LLMs’ strategic reasoning through agent interaction and prompting (Gandhi et al., 2023; Fu et al., 2023; Xia et al., 2024; Hua et al., 2024), though these efforts have primarily focused on simpler single-issue settings. In contrast, MIB scenarios demand more complex reasoning and adaptive strategies. Here, LLMs often produce irrational or suboptimal offers, falling short not only of human-level reasoning but also of achieving the core negotiation goal of outcome maximization (Schneider et al., 2023; Kwon et al., 2024).

LLMs’ Action with Tooling Recent work has integrated external tools, such as APIs and function calls, to enhance LLMs’ reasoning and decision-making (Zhang et al., 2024b; Yao et al., 2022; Qin et al., 2023). Building on this, we pioneer the use of tooling in complex interactions like negotiation. By integrating a linear programming solver as a dynamic reward function for turn-level offer optimization, we enhance LLM-based agents’ strategic reasoning and adaptability. This advances research on strategic decision-making in complex interactive settings, including multi-issue bargaining.

7 Conclusion

We proposed *ASTRA*, an adaptive and strategic framework for turn-level offer optimization, enabling agents to maximize its outcome. In the

challenging MIB scenario, *ASTRA* outperformed baselines across diverse partner agents and was validated through human evaluation as highly strategic and effective in achieving negotiation goals. Dynamic action with LP in *ASTRA* achieves the adaptability of a reward-based framework without training costs while ensuring Pareto-optimal offers, providing an efficient and scalable solution for strategic negotiation. *ASTRA*’s turn-level strategies, optimal offers, and opponent behavior signals enhance human decision-making in negotiation-assistive systems. Future work will explore its integration into coaching and decision-support applications, where *ASTRA* may serve not only as a negotiation engine but also as a powerful tool for guiding and improving human strategic behavior.

8 Broader Impact and Ethical Considerations

8.1 Datasets

Our study uses a publicly available, anonymized negotiation dataset (Chawla et al., 2021), ensuring compliance with its licensing terms, intended use, and ethical guidelines. The dataset is in English, and all simulations between the baseline and our agent were conducted in English. Since negotiation strategies can vary across cultures (Luo, 2008), our findings may not directly generalize to interactions conducted in other languages. We encourage future research to explore negotiation dynamics in multilingual settings, analyzing how agents trained in different languages interact and make strategic decisions. Our study serves as a foundational step in this direction, providing insights that can inform further investigations into culturally adaptive negotiation strategies.

8.2 LLMs

Our use of LLMs adhered strictly to their intended purpose and licensing terms, aligning with ethical and regulatory standards. Similar to recent studies conducting LLM-based agent simulations, our approach ensures responsible and transparent evaluation while maintaining compliance with operational guidelines.

8.3 Human Evaluation

To assess the “Strategicness” of responses generated by our models, we first relied on six expert annotators with domain expertise from our research institute. In addition, we conducted a complementary human evaluation to assess ASTRA’s potential as a decision-support tool. All evaluations followed a consistent and blinded protocol to ensure fairness and objectivity. Given the complexity of negotiation dialogue evaluation, we prioritized expert judgment over general crowdsourcing to ensure a more rigorous and informed assessment. Prior to participation, annotators were fully informed about the purpose of the evaluation and its potential impact, and their consent was obtained.

8.4 Use AI assistant Tools

We used AI-based tools such as ChatGPT for language refinement and assisting with code debugging and optimization. However, all conceptual contributions, experimental designs, core algorithms, methodology, and final implementations

were independently developed, reviewed, and validated by the authors.

9 Limitations

Our current work demonstrates the effectiveness of our ASTRA framework in a single Multi-Issue Bargaining (MIB) scenario involving three camping-related items. Through enhanced adaptability and strategic reasoning, our agent successfully navigates negotiation dynamics within this setting. However, this serves as a proof of concept, and the framework can be scalably extended to other negotiation and game-theoretic scenarios that can be formulated using LP with more complex issue spaces.

Additionally, due to the inherent limitations of LLMs in reasoning under imperfect information and complex interaction dynamics, our approach currently optimizes LP-based offer generation by adjusting specific LP parameters. As LLMs continue to advance in reasoning capabilities, future work can explore fully dynamic LP problem formulation within negotiation contexts, where the LLM autonomously codes and adapts optimization models in real-time.

Furthermore, while our human evaluation was conducted at the turn level, assessing how effectively ASTRA-generated offers contribute to strategic goal achievement within a given negotiation context, a crucial next step is to validate our agent’s performance in direct interactions with human negotiators. This would provide deeper insights into its real-world applicability. Moreover, as ASTRA also functions as a negotiation coaching tool, future work can assess how well it enhances human negotiators’ performance by providing actionable feedback and optimized offer suggestions throughout the negotiation process.

10 Acknowledgments

Research was sponsored by the Army Research Office under Cooperative Agreement Number W911NF-25-2-0040. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes, notwithstanding any copyright notation herein.

References

Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. 2024. Cooperation, competition, and maliciousness: Llm-stakeholders interactive negotiation. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Robert Axelrod and William D Hamilton. 1981. The evolution of cooperation. *science*, 211(4489):1390–1396.

Tim Baarslag, Mark JC 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*, 30:849–898.

Tim Baarslag, Koen Hindriks, and Catholijn Jonker. 2013. A tit for tat negotiation strategy for real-time bilateral negotiations. *Complex automated negotiations: theories, models, and software competitions*, pages 229–233.

Yoella Bereby-Meyer and Muriel Niederle. 2005. Fairness in bargaining. *Journal of Economic Behavior & Organization*, 56(2):173–186.

Federico Bianchi, Patrick John Chia, Mert Yuksekgonul, Jacopo Tagliabue, Dan Jurafsky, and James Zou. 2024. How well can llms negotiate? negotiationarena platform and analysis. *arXiv preprint arXiv:2402.05863*.

Jakub Brzostowski and Ryszard Kowalczyk. 2006. Adaptive negotiation with on-line prediction of opponent behaviour in agent-based negotiations. In *2006 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 263–269. IEEE.

Kushal Chawla, Gale Lucas, Jonathan May, and Jonathan Gratch. 2022. Opponent modeling in negotiation dialogues by related data adaptation. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 661–674, Seattle, United States. Association for Computational Linguistics.

Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May, and Jonathan Gratch. 2021. Casino: A corpus of campsite negotiation dialogues for automatic negotiation systems. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

Kushal Chawla, Ian Wu, Yu Rong, Gale Lucas, and Jonathan Gratch. 2023. Be selfish, but wisely: Investigating the impact of agent personality in mixed-motive human-agent interactions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13078–13092, Singapore. Association for Computational Linguistics.

Jiangjie Chen, Siyu Yuan, Rong Ye, Bodhisattwa Prasad Majumder, and Kyle Richardson. 2023. Put your money where your mouth is: Evaluating strategic planning and execution of llm agents in an auction arena. *arXiv preprint arXiv:2310.05746*.

Yang Deng, Lizi Liao, Liang Chen, Hongru Wang, Wenqiang Lei, and Tat-Seng Chua. 2023. Prompting and evaluating large language models for proactive dialogues: Clarification, target-guided, and non-collaboration. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10602–10621, Singapore. Association for Computational Linguistics.

Ernst Fehr and Klaus M Schmidt. 1999. A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868.

Chaim Fershtman. 1990. The importance of the agenda in bargaining. *Games and Economic Behavior*, 2(3):224–238.

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

Kanishk Gandhi, Dorsa Sadigh, and Noah D Goodman. 2023. Strategic reasoning with language models. *arXiv preprint arXiv:2305.19165*.

Werner Güth, Rolf Schmittberger, and Bernd Schwarze. 1982. An experimental analysis of ultimatum bargaining. *Journal of economic behavior & organization*, 3(4):367–388.

He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. Decoupling strategy and generation in negotiation dialogues. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2333–2343, Brussels, Belgium. Association for Computational Linguistics.

Henrike Heunis, Niels J Pulles, Ellen Giebels, Bas Kolöffel, and Aldis G Sigurdardottir. 2024. Strategic adaptability in negotiation: a framework to distinguish strategic adaptable behaviors. *International Journal of Conflict Management*, 35(2):245–269.

Harold M Hochman and James D Rodgers. 1969. Pareto optimal redistribution. *The American economic review*, 59(4):542–557.

Wenyue Hua, Ollie Liu, Lingyao Li, Alfonso Amayuelas, Julie Chen, Lucas Jiang, Mingyu Jin, Lizhou Fan, Fei Sun, William Wang, et al. 2024. Game-theoretic llm: Agent workflow for negotiation games. *arXiv preprint arXiv:2411.05990*.

Yin Jou Huang and Rafik Hadfi. 2024. How personality traits influence negotiation outcomes? a simulation based on large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10336–10351, Miami, Florida, USA. Association for Computational Linguistics.

Emmanuel Johnson. 2019. Using automated agents to teach negotiation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9888–9889.

Catholijn Jonker, Reyhan Aydogan, Tim Baarslag, Katsuhide Fujita, Takayuki Ito, and Koen Hindriks. 2017. Automated negotiating agents competition (anac). In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.

Harold H Kelley. 1996. *A classroom study of the dilemmas in interpersonal negotiations*. Berkeley Institute of International Studies.

Deuksin Kwon, Emily Weiss, Tara Kulshrestha, Kushal Chawla, Gale Lucas, and Jonathan Gratch. 2024. *Are LLMs effective negotiators? systematic evaluation of the multifaceted capabilities of LLMs in negotiation dialogues*. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 5391–5413, Miami, Florida, USA. Association for Computational Linguistics.

Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. 2023. Reward design with language models. *arXiv preprint arXiv:2303.00001*.

Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017a. *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.

Mike Lewis, Denis Yarats, Yann N Dauphin, Devi Parikh, and Dhruv Batra. 2017b. Deal or no deal? end-to-end learning for negotiation dialogues. *arXiv preprint arXiv:1706.05125*.

Peng Luo. 2008. Analysis of cultural differences between west and east in international business negotiation. *International Journal of Business and Management*, 3(11):103–106.

M Olekalns and LR Weingart. 2008. Emergent negotiations: Stability and shifts in negotiation dynamics. *negotiation and conflict management research*, 1 (2), 135–160.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. 2023. Toollm: Facilitating large language models to master 16000+ real-world apis, 2023. URL: <https://doi.org/10.48550/ARXIV/2307.06347>.

Johannes Schneider, Steffi Haag, and Leona Chandra Kruse. 2023. Negotiating with llms: Prompt hacks, skill gaps, and reasoning deficits. *arXiv preprint arXiv:2312.03720*.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Ryan Shea, Aymen Kallala, Xin Lucy Liu, Michael W Morris, and Zhou Yu. 2024. Ace: A llm-based negotiation coaching system. *arXiv preprint arXiv:2410.01555*.

Herbert A Simon. 1990. Bounded rationality. *Utility and probability*, pages 15–18.

Nancy A Welsh. 2003. Perceptions of fairness in negotiation. *Marq. L. Rev.*, 87:753.

Tian Xia, Zhiwei He, Tong Ren, Yibo Miao, Zhuosheng Zhang, Yang Yang, and Rui Wang. 2024. *Measuring bargaining abilities of LLMs: A benchmark and a buyer-enhancement method*. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 3579–3602, Bangkok, Thailand. Association for Computational Linguistics.

Runzhe Yang, Jingxiao Chen, and Karthik Narasimhan. 2021. *Improving dialog systems for negotiation with personality modeling*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 681–693, Online. Association for Computational Linguistics.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.

Yadong Zhang, Shaoguang Mao, Tao Ge, Xun Wang, Adrian de Wynter, Yan Xia, Wenshan Wu, Ting Song, Man Lan, and Furu Wei. 2024a. Llm as a mastermind: A survey of strategic reasoning with large language models. *arXiv preprint arXiv:2404.01230*.

Yinger Zhang, Hui Cai, Xierui Song, Yicheng Chen, Rui Sun, and Jing Zheng. 2024b. *Reverse chain: A generic-rule for LLMs to master multi-API planning*. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 302–325, Mexico City, Mexico. Association for Computational Linguistics.

Zheng Zhang, Lizi Liao, Xiaoyan Zhu, Tat-Seng Chua, Zitao Liu, Yan Huang, and Minlie Huang. 2020. *Learning goal-oriented dialogue policy with opposite agent awareness*. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 122–132, Suzhou, China. Association for Computational Linguistics.

A Components in Negotiation dialogue System

A.1 Partner Preference Consistency Checker

In negotiation, a player’s score is calculated as the inner product of item quantities obtained and their preferences. The agent predicts the partner’s score using inferred partner preferences (IPP), which are critical for generating strategic offers. A Consistency Checker ensures the IPP aligns with the partner’s behavior, identifying inconsistencies in two cases:

- 1) when the partner’s score from their offer at time t ($S_p^{P,(t)}$) is lower than their score from the agent’s offer at time $t-1$ ($S_p^{A,(t-1)}$), indicating $S_p^{P,(t)} < S_p^{A,(t-1)}$).
- 2) when the partner’s utterances explicitly contradict the current IPP, such as expressing $V_{\text{food}} > V_{\text{water}}$ while IPP suggests $V_{\text{food}} < V_{\text{water}}$.

Upon detecting inconsistency, the agent recalculates the IPP using updated information from the partner’s utterances and offers, ensuring accurate predictions and strategic decision-making.

A.2 First Stage in ASTRA

In determining fairness (see Equation 1), we rely on insights from prior literature. Perceptions of unfairness in reward distribution generally arise when an individual’s share falls below 20–30% of the total. The inequity aversion model proposed by [Fehr and Schmidt \(1999\)](#) specifies a theoretical acceptance threshold given by $s = \frac{\alpha}{1+2\alpha}$, where α , typically ranging from 0.3 to 0.75, captures sensitivity to disadvantageous inequality. This results in a predicted cutoff between 15% and 28% of the total allocation. Empirical findings from ultimatum game studies ([Güth et al., 1982](#); [Bereby-Meyer and Niederle, 2005](#)) are consistent with this prediction, showing that offers below 20% are rejected in approximately 50% of cases.

To accommodate individual variation in inequity sensitivity, we conservatively set the fairness threshold at 35% of the total score (i.e., $\theta_f = 4$, corresponding to a 4-point difference from the counterpart), which clearly exceeds both theoretical and empirical thresholds and constitutes a practically acceptable division.

A.3 Second Stage in ASTRA

In the second stage of ASTRA, the agent uses an LP solver to generate optimal offers based on the nego-

tiation context. Figure 8 shows the Python function used for turn-level offer optimization, incorporating dynamically adjusted parameters derived from the partner’s behavioral signals.

A.4 Third Stage in ASTRA

For the *PAP* in the third stage, we assign weights to *TS* and *SI* to compute the *PAP* score. To determine an appropriate weighting scheme, we implemented separate Virtual Partner Agents (VPAs) within ASTRA for each weight combination and ran 50 agent-to-agent simulations per VPA setting. Given that the final score (*TS*) has a stronger impact on the partner agent’s decisions, we prioritized it by testing higher weights, specifically 0.5, 0.75, and 1.0, relative to *SI*. Based on the experimental results, we selected weights of 0.75 (w) for *TS* and 0.25 for *SI*.

TS weight (w)	SI weight ($1-w$)	Avg. Score-All (P1 vs P2)	Avg. Score-Agree. (P1 vs P2)
0.5	0.5	19.71 vs 13.04	21.67 vs 14.11
0.75	0.25	20.69 vs 12.33	24.34 vs 14.50
1.0	0.0	20.51 vs 12.8	24.03 vs 15

Table 5 illustrates various tactics used in the third stage for strategic offer assessment, along with detailed descriptions specifying what, when, and how each tactic is applied.

B Agent simulations

B.1 Model Hyperparameters

For consistency, the default decoding settings were used for each LLM model in the simulation. All models were run with temperature = 1.0; only top- p differed:

- GPT-4o (mini): top- p = 1.0
- Claude-3.5-Sonnet: top- p = 0.99
- Gemini-2.0-Flash: top- p = 0.95

B.2 Simulation Examples

Table 7 shows an example of a simulation between our ASTRA-powered agent and a partner agent in a standard multi-issue bargaining scenario, demonstrating ASTRA’s ability to strategically adapt its offers. Table 8 extends this to a more complex setting with four heterogeneous issue types—integer, categorical, and binary—highlighting the framework’s generalizability and robustness with minimal modification. Finally, Table 9 presents representative errors from the partner agent, illustrating the limitations of LLM-based reasoning when utility structure and constraints are not explicitly modeled.

C Dataset

We utilize the CaSiNo dataset to set up agent simulation scenarios and to train our RL negotiation agent as a baseline. As shown in the Table 6, CaSiNo is a corpus consisting of 1,030 dialogues specifically designed for automated negotiation research. The dialogues involve two participants acting as campsite neighbors negotiating food, water, and firewood packages for a camping trip. Each item (food, water, firewood) is randomly assigned a priority (high, medium, low) per participant. Negotiators aim to maximize their individual points, with high, medium, and low priority items worth 5, 4, and 3 points, respectively. The dataset also includes annotations for negotiation strategies, categorized as either prosocial behaviors (e.g., small-talk, empathy) or proselytizing behaviors (e.g., undervaluing the partner, emphasizing one’s own needs). However, we did not use these strategy annotations when training our RL baseline model. Detailed examples from the dataset can be found in (Chawla et al., 2021).

D Baselines

We implement the baseline models following the original methods proposed in their respective papers. For the RL agent in particular, we leverage state-of-the-art RL algorithms to build a stronger and more competitive baseline. The implementation details of each baseline are described below.

D.1 RL Agent

We adopt the approach from Chawla et al. (2023), which refines Lewis et al. (2017a)’s method by training a selfish RL agent using a reward function centered on self-interest. To enhance language capabilities, we fine-tune the pre-trained Flan-T5 model² on the CaSiNo dataset instead of training an Recurrent Neural Network (RNN) from scratch. We also replace their REINFORCE algorithm with the more advanced Proximal Policy Optimization (PPO) (Schulman et al., 2017) for improved optimization.

We train our RL agent using a two-stage self-play framework. In Stage 1, we fine-tune the model (i.e., Flan-T5-Base (2.5M)) in a supervised manner on the Casino dataset with an 8/1/1 train/validation/test split. The model is trained for 10 epochs with a learning rate of 5e-5, using a lin-

ear learning rate scheduler and a batch size of 16. Optimization is performed with AdamW.

In Stage 2, we refine the agent using PPO with self-play, where the fine-tuned model iterates across diverse agent-partner preference combinations per epoch. The policy is trained for 15 epochs with a batch size of 8 and a clip ratio of 0.2. All models are trained on A100 GPUs.

D.2 Pro-CoT

Following the approach by Deng et al. (2023), we prompt the negotiation agent (based on GPT-4o) to generate descriptive thoughts outlining its intermediate reasoning and planning steps for achieving negotiation goals (e.g., maximizing its score), using a Chain-of-Thought prompting scheme. In particular, the agent is guided to proactively and strategically lead the dialogue by generating an appropriate strategy and dialogue action in the given context.

D.3 ICL-AIF

Following Fu et al. (2023), we introduce an AI feedback loop where, at the end of each negotiation session, a Critic model (based on GPT-4o) provides feedback on the previous dialogue to improve the agent’s performance. The negotiation agent incorporates both the prior negotiation history and the received feedback as in-context demonstrations, and then proceeds to negotiate again under the same setting. After each session, the Critic suggests three strategies, and this process is repeated for three iterations. The highest score achieved across the three sessions is considered the agent’s final performance.

E Abaltion Study

Figure 9 visualizes the impact of adjusting weights (i.e., α and β) between Partner Acceptance Probability (PAP) and Strategy Assessment (SA) Score in ASTRA’s third stage, based on agent simulation results.

F Human Evaluation

To evaluate the effectiveness of the agent and proposed framework, expert annotators with negotiation knowledge were each given a set of questionnaires via Qualtrics. An example of the questionnaire can be found in Figure 10 and 11.

²<https://huggingface.co/google/flan-t5-base>

Function of Linear Programming Solver

```
def solve_LP(max_point, lambda, agents_value, partner_value):
    """
    Solves a Linear Programming problem for the given parameters.
    Args:
    max_point (int): Maximum points the agent can get.
    lambda (float): Lambda balancing both parties' objectives. [0-1]
    agents_value (dict): Agent's values for food, water, and firewood.
    partner_value (dict): Partner's values for food, water, and firewood.
    Returns:
    tuple: A tuple containing the calculated scores and item allocations.
    """
    import pulp

    A_vals, B_vals = agents_value, partner_value
    A_F, A_W, A_FW = A_vals['food'], A_vals['water'], A_vals['firewood']
    B_F, B_W, B_FW = B_vals['food'], B_vals['water'], B_vals['firewood']

    # Define the LP problem
    problem = pulp.LpProblem("Maximize_Points", pulp.LpMaximize)

    # Define variables
    X = pulp.LpVariable("X", 0, 3, cat='Integer') # Food agent gets
    Y = pulp.LpVariable("Y", 0, 3, cat='Integer') # Water agent gets
    Z = pulp.LpVariable("Z", 0, 3, cat='Integer') # Firewood agent gets

    # Objective function
    objective = (
        (A_F * X + A_W * Y + A_FW * Z) +
        (1 - lambda) * ((B_F * (3 - X) + B_W * (3 - Y) + B_FW * (3 - Z)))
    )
    problem += objective

    # Constraints
    problem += A_F * X + A_W * Y + A_FW * Z <= max_point,
    problem += A_F * X + A_W * Y + A_FW * Z >= 10,
    problem += B_F * (3 - X) + B_W * (3 - Y) + B_FW * (3 - Z) >= 5,

    # Solve the problem
    problem.solve()

    return X.varValue, Y.varValue, Z.varValue # Solution
```

Figure 8: Python function for LP solver execution in ASTRA

Turn-level Collaborative Tactics			
Initial Concession (LIC)	When	Early in the negotiation; no initial concessions have been made.	
	How	Make a significant concession, giving the partner something they highly value.	
	Why	Signals cooperation and encourages reciprocal concessions.	
Continued-Smaller Concessions (CSC)	When	After a large concession, if the partner responds with a small or no concession.	
	How	Make a smaller concession	
	Why	Maintains cooperation and encourages continued reciprocity.	
Reciprocal Concessions (RC)	When	The partner makes a significant concession	
	How	Respond with a concession	
	Why	Rewards cooperation and promotes fair exchanges.	
Logrolling (LGR)	When	Anytime in the negotiation.	
	How	Trade items based on relative importance - give what they value less in exchange for what you value more.	
	Why	Maximizes joint gains and fosters win-win outcomes.	
Mutual Gain Focus (MGF)	When	The partner's latest offer gives them a higher score than yours.	
	How	Adjust the offer to improve mutual benefits, emphasizing fairness.	
	Why	Builds cooperation and encourages a collaborative environment.	
Turn-level Competitive Tactics			
Aggressive Early Offers (AEO)	When	At the start of negotiation, before many offers are exchanged.	
	How	Make highly self-favorable offers	
	Why	Establishes a strong position, allowing room for later concessions.	
Response to Extreme Offer (REO)	When	The partner makes an extremely selfish offer	
	How	Counter with a highly self-favorable offer, matching their approach.	
	Why	Discourages extreme demands and promotes fairness.	
No Concession Response (NCR)	When	The partner remains firm, making no concessions in their last two offers.	
	How	Stop conceding and hold your position	
	Why	Prevents exploitation and maintains negotiation balance.	
Reject Negative Concession (RNC)	When	The partner's last two offers increased their own score	
	How	Maintain your position without making further concessions.	
	Why	Signals resistance to unfair demands and encourages balance.	

Table 5: Description of various turn-level tactics used in the strategic assessment of ASTRA 3rd stage. Useful negotiation tactics (Kelley, 1996; Baarslag et al., 2016) are operationalized at the turn level for offer proposals.

Preference & Arguments	
Speaker	Conversation
P1	[High] Water: We like to go on runs and it increases the need of this. [Medium] Food: Food overall is a good mood booster. [Low] Firewood: We do not care for fire and it is not necessary to us.
P2	[High] Food: Food really increases everyone's morale. [Medium] Firewood: We like to have a large fire. [Low] Water: We don't drink water that often.
P1	How are you today? Did you have any preferences on the supplies we will be trading?
P2	I am good. How about yourself? I think I would like some firewood to start off with. We like to have bigger fires. What about you?
P1	I am good as well.
P2	Water is a little important to us too though, if possible maybe we can split that or maybe we can get some more food in replacement.
P1	That may be possible.... What did you have in mind for the food replacement?
P2	You can have all the water if we can have all the food?
P1	I dont think I am okay with that. Food is essential to our groups morale when camping. We would like 1 additional food preferably.
P2	Well you guys did say you did not care much about large fires. What if you gave all the firewood in replace for the water and you can still keep 1 food?
P1	So I would get 3 water and 1 food and youd get 3 firewood and 2 food?
P2	Yea that seems like an alright trade to me.
P1	Hmm... alright then
P2	Submit-Deal

Table 6: An example of negotiation setting and dialogue from Chawla et al. (2021)

Speaker	Conversation	Score		Partner Behavior		Selected strategy
		Agent	Partner	Joint	Fairness	
[Integrative Case] NegoAgent priorities: Food=High (5), Water=Middle (4), Firewood=Low (3) PartnerAgent priorities: Food=Low (3), Water=Middle (4), Firewood=High (5)						
PartnerAgent	Hello! let's start the negotiation!					
NegoAgent	Glad to kick off our negotiation! You mentioned your priorities earlier, which is great to know. To clarify, could you tell me what your highest priority item is? (Asking Preference)					
PartnerAgent	I appreciate your interest in understanding my priorities. My highest priority is firewood. I'd love to hear about your priorities as well so we can find a beneficial agreement.					
NegoAgent	Thanks for sharing your highest priority! It's interesting that firewood is at the top for you. Just to clarify, is food your least favorite item in this negotiation? (Asking Preference)					
PartnerAgent	Yes, food is my least preferred item in this negotiation. Could you tell me more about your priorities? That might help us reach a mutually beneficial agreement.					
NegoAgent	Considering your priority on firewood, I suggest that we agree on an offer where I take all 3 food, all 3 water packages, and 1 firewood package. This way, you'll still receive 2 firewood packages, which should help satisfy your needs, while I ensure we have ample food and water for our trip. What do you think? (Offer)	30	10	40	Unk.	Unk. AIO
PartnerAgent	I understand where you're coming from. However, I would still need some water for balance in our supplies. How about I take 2 firewood, 2 water, and 1 food, while you take 2 food, 1 water and 1 firewood? This way, we both get what we value.	17	21	38	Unfair	Neu.
NegoAgent	I appreciate your suggestions so far, and I understand your needs for water and firewood. I propose that I take all 3 food packages, all 3 water packages, and 1 firewood package. This way, you will still receive 2 firewood packages, which aligns with your highest priority. I believe this arrangement provides a good balance of supplies for both of us. What do you think? (Counter Offer)	30	10	40		AEO
PartnerAgent	I understand your position well, and I truly value a fair deal. Given your proposal, how about I take 2 firewood, 2 water, and 1 food while you take the 3 food items and 1 firewood? This would give you the majority of the food you seek and some firewood while ensuring I have a fair share of the items important to me, particularly the firewood and water. Let me know what you think!	17	21	38	Unfair	Neu.
NegoAgent	I appreciate your willingness to negotiate, and I think we can reach a great agreement. I propose that I take all 3 food packages, all 3 water packages, and 1 firewood package. This way, you would receive 2 firewood packages. I believe this allocation balances our needs and ensures we have a good supply for the trip. Let me know your thoughts! (Counter Offer)	30	10	40		NCR
PartnerAgent	I appreciate your clear preferences for food and water. To find a common ground, I propose that I take 2 firewood, 1 water, and 2 food, while you take 1 firewood, 2 food, and 2 water. This adjustment gives you more water while I maintain my priority on firewood. Let me know if this works for you.	16	20	36	Unfair	Gen.
NegoAgent	Thank you for your continued collaboration in this negotiation. Given our discussions and understanding of each other's priorities, I propose that I take all 3 food packages, 2 water packages, and 1 firewood package. This means you would receive 2 firewood and 1 water package, which should still support your needs for firewood while allowing me to gather adequate food and water for our trip. I believe this arrangement strikes a fair balance. What are your thoughts? (Counter Offer)	26	14	40		RC
[ASTRA: Offer selection] » Fairness=unfair, Stance=generous LP parameters: max_bound=30, lambda=0.3 » Selected Strategy: RC » Offer candidates: Offer (0): (30: f=3 w=3 fw=1) Final Score: 0.10 PAP: 0.3 (TS: 0.22 SI: 0.5) SA: 0.0 Offer (1): (27: f=3 w=3 fw=0) Final Score: 0.66 PAP: 0.5 (TS: 0.46 SI: 0.6) SA: 0.75 Offer (2): (26: f=3 w=2 fw=1) Final Score: 0.85 PAP: 0.56 (TS: 0.52 SI: 0.64) SA: 1.0 Offer (3): (23: f=3 w=2 fw=0) Final Score: 0.59 PAP: 0.75 (TS: 0.76 SI: 0.74) SA: 0.5 Offer (4): (22: f=3 w=1 fw=1) Final Score: 0.39 PAP: 0.64 (TS: 0.6 SI: 0.72) SA: 0.25 » **Finally Selected Best Offer** : (26: f=3 w=2 fw=1)						
PartnerAgent	Based on your latest proposal, I understand your higher priority for food and water. I propose the following: I receive 2 firewood packages, 2 water packages, and 0 food, while you receive all 3 food packages, 1 water package, and 1 firewood package. This way, I maintain my priority on firewood and water while you get the food packages you value. Let me know if this works for you!	22	18	40	Fair	Gen.
NegoAgent	Thank you for your proposals so far. Based on our discussions, I would like to suggest that I take all 3 food packages and 2 water packages , while you receive 0 food packages, 1 water package, and 3 firewood packages. I believe this arrangement supports my need for adequate food and water while still respecting your priority on firewood. Let me know your thoughts on this proposal! (Counter Offer)	23	19	42		RC
PartnerAgent	ACCEPT-DEAL					

Table 7: Example simulation dialogue between the NegoAgent with ASTRA framework and Partner Agent. Red brackets indicate the response mode for each turn, while the blue text represents the corresponding utterance. Additionally, in ASTRA, the inferred behavioral signals of the Partner Agent (i.e., Fairness and Stance) and the final selected strategy for adaptive decision-making are provided. Here, Unk. stands for Unknown, Neu. for Neutral, and Gen. for Generous.

Scenario Description										
Scenario: Joint Research Resource Allocation										
Description: Two researchers are negotiating the allocation of limited resources for a joint research project. The available resources consist of 5 units of experimental equipment, 5 research staff members, a shared lab environment (computer, chemistry, or biology), and the possibility of using the lab over the weekend.										
Each participant has different priorities regarding these resources. The agent and the partner each assign distinct levels of importance to four key issues: 1. Equipment (E) (integer, 0-5) : Number of experimental tools allocated to each party. 2. Staff (S) (integer, 0-5) : Number of data collection personnel assigned. 3. Lab Type (L) (categorical) : Shared lab environment chosen from computer, chemistry, or biology. 4. Weekend Use (W) (binary) : Whether the lab is available for weekend use (True or False).										
<pre>Variable Adjustment in LP Solver def calculateBestOfferFromLP(maximum_value, lambda_value, agents_value, partner_value): ... problem = pulp.LpProblem("Research_Allocation", pulp.LpMaximize) # Integer Issue E = pulp.LpVariable("Equipment", 0, 5, cat='Integer') S = pulp.LpVariable("Staff", 0, 5, cat='Integer') # Categorical Issue L_computer = pulp.LpVariable("L_computer", 0, 1, cat='Binary') L_chemistry = pulp.LpVariable("L_chemistry", 0, 1, cat='Binary') L_biology = pulp.LpVariable("L_biology", 0, 1, cat='Binary') problem += L_computer + L_chemistry + L_biology == 1 # Binary Issue W = pulp.LpVariable("WeekendUse", 0, 1, cat='Binary') ... </pre>										
Issue Priority (value)										
NegoAgent	[High (4)]: Staff									
	[Medium-High (3)]: Equipment									
PartnerAgent	[Medium-Low (2)]: Lab (High (1)=Biology, Medium (0.6)=Chemistry, Low (0.2)=Computer)									
	[Low (1)]: Weekend									
Context										
Speaker		Utterance	Partner Score	Agent Score						
PartnerAgent	...		26	12.4						
	I'd like to take 4 pieces of equipment, 3 staff, use the computer lab, and include weekend access. The computer lab is more suitable for my work, and I really need weekend access.									
NegoAgent	You're taking too much equipment. I value research staff more, so I will take 4 staffs. I'm okay with weekend access.		17.2	25						
	But I need at least 2 units of equipment. Let's switch to the biology lab instead.									
PartnerAgent	Biology gives me the least value, so how about we compromise with the chemistry lab? I'll reduce to 3 equipment and 3 staff.		21.6	16.2						
	Well... how about I take 1 piece of equipment and just 4 staff members, and we use the biology lab?									
[ASTRA: Offer selection]										
» Fairness=unfair, Stance=generous LP parameters: max_bound=23, lambda=0.7										
» Selected Strategy: RC										
» Offer candidates:										
Offer (1): (23: E=4 S=2 L=biology W=True)										
Offer (2): (23: E=0 S=5 L=biology W=True)										
Offer (3): (22: E=5 S=1 L=biology W=True)										
Offer (4): (22: E=1 S=4 L=biology W=True)										
Offer (5): (21: E=0 S=5 L=computer W=True)										
Offer (6): (21: E=2 S=3 L=biology W=True)										
Offer (7): (20: E=1 S=4 L=computer W=True)										
» Finally Selected Best Offer :										
(22: E=1 S=4 L=biology W=True) Final Score: 0.89 PAP: 0.69 (TS: 0.72 SI: 0.58) SA: 1										

Table 8: Demonstration of ASTRA in an extended negotiation scenario. This scenario involves four issues with diverse types (integer, categorical, binary), exceeding the complexity of CaSiNo (Chawla et al., 2021). With minimal adaptations to the LP solver, ASTRA remains effective and generalizable.

Speaker	Conversation	Error type
PartnerAgent	<p>..</p> <p>I appreciate your consideration. Here's a revised offer that might work for both of us: I will take 3 firewood, 1 food, and 1 water, while you will take 3 food and 2 water.</p> <p>This way, you get all the food and most of the water, and I maintain my priority for firewood. Let me know your thoughts on this proposal.</p> <p>...</p>	Item Mis-Specification (Over-Specified)
PartnerAgent	Hello! let's start the negotiation!	Preference
NegoAgent	<p>Hello! It's great to start the negotiation. I appreciate your readiness to jump right into things.</p> <p>Since we haven't discussed priorities yet, could you let me know what your highest priority item is?</p> <p>I'll take 3 firewood, 2 water, and 1 food.</p>	Inconsistency
PartnerAgent	<p>This gives me the maximum from my highest valued item, firewood.</p> <p>Could you let me know what your priorities are so we can ensure both of our needs are met?</p>	
NegoAgent	<p>Thank you for sharing your initial item choices! I see that you've prioritized firewood as your highest valued item.</p> <p>To finalize our discussion, could you clarify if food is your lowest priority?</p>	
PartnerAgent	<p>Yes, food is my lowest priority because it provides the least value to me.</p> <p>Could you tell me your top priority so we can try to find a middle ground?</p> <p>...</p>	
NegoAgent	<p>Taking into account our needs, I propose that I take 3 food packages, 2 water packages, and no firewood. This way, I will have enough food and water for our trip.</p> <p>Consequently, you would receive 0 food, 1 water, and 3 firewood packages. (Score: 19)</p> <p>I believe this reflects a fair distribution given the situation. What are your thoughts on this?</p>	
PartnerAgent	<p>I propose that I take 2 firewood, 1 water, and 1 food package. (Score: 17)</p> <p>You would then have 3 food, 2 water, and 1 firewood package.</p>	
[Consistency Checker]	<p>This will ensure a balanced distribution based on both of our priorities. How does this sound?</p> <p>Inconsistency detected: partner score in agent offer (19) > partner score in partner offer (17)</p> <p>Thank you for your continued engagement in this discussion!</p>	
NegoAgent	<p>From our chat, it seems like your priorities are different from what I thought.</p> <p>Could you clarify what your highest priority item is?</p> <p>Is it firewood, or do you have another item in mind that you value more?</p> <p>...</p>	

Table 9: Error cases from the PartnerAgent in the agent-to-agent simulations. Here, NegoAgent refers to our agent utilizing ASTRA.

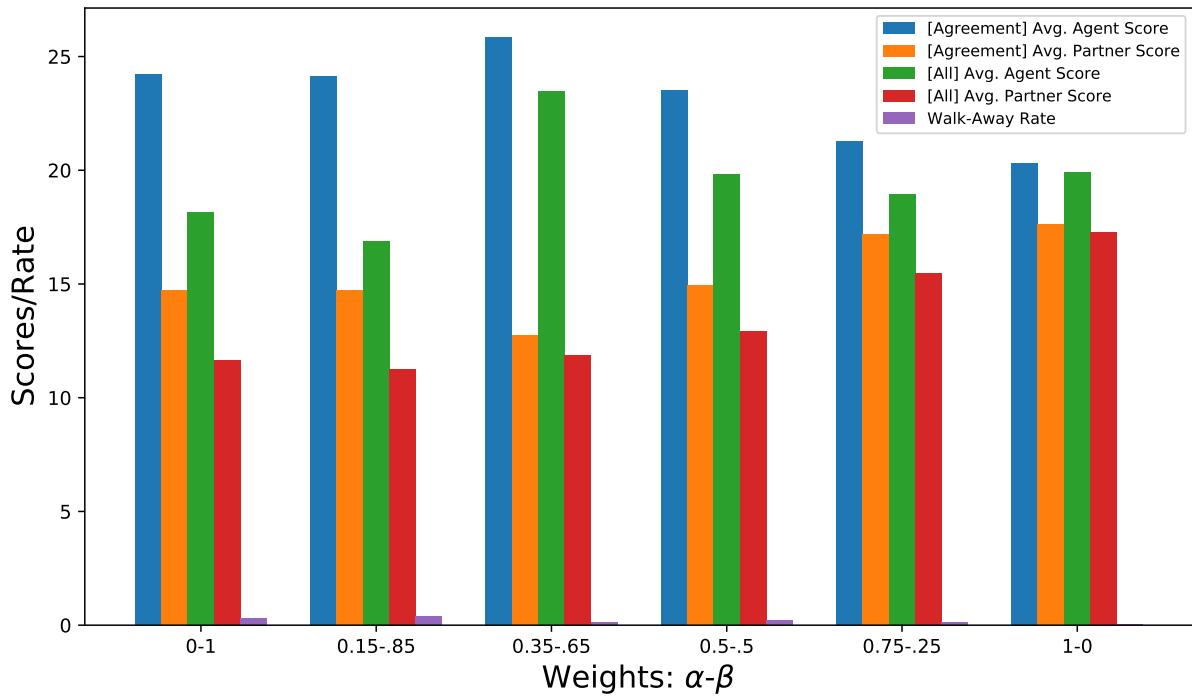


Figure 9: Score and Walk-Away ratio with Weight variations

Q3

Q3. Evaluate how strategic the Agent's final response is in achieving its negotiation goals, based on the given both participant's preferences for items and the Partner, as well as the conversation history.

Players' preferences for the 3 items:
 (High priority: red, Low priority: blue)
 Agent: food=high(5), water=middle(4), firewood=low(3)
 Partner: food=high(5), water=middle(4), firewood=low(3)

Conversation history:
 [PARTNER]: Hello! let's start the negotiation!
 [AGENT]: I see you're ready to start the negotiation! It's great to have that setup. Since we haven't delved into the item priorities yet, can I ask you, what is your highest priority item?
 [PARTNER]: My highest priority item is food because it's crucial for our camping trip. How about you? What is your highest priority item?
 [AGENT]: I appreciate you sharing that your highest priority item is food; it's certainly vital for a camping trip! Now, I'd like to clarify something regarding your priorities. What is your lowest priority item?
 [PARTNER]: My lowest priority is firewood. How about our first offer: I take 3 food, 1 water, and 0 firewood, and you take 2 water and 3 firewood. What do you think?
 [AGENT]: I appreciate your offer, but I think we can find a more balanced agreement. I propose that I take 3 food, 3 water, and 1 firewood, while you will receive 0 food, 0 water, and 2 firewood. This way, we ensure a better distribution for our camping needs. What do you think?
 [PARTNER]: I understand your desire for a better distribution. How about I take 2 food, 1 water, and 1 firewood, while you take 1 food, 2 water, and 2 firewood? This way, we each get a fair amount of high-priority items.

Agent's response: I appreciate our ongoing discussions, and I believe we can reach a balanced agreement. With that in mind, I propose that I take 2 food, 3 water, and 2 firewood, while you receive 1 food, 0 water, and 1 firewood. This way, we both have a fair share of what we need for our trip.

Not very strategic
 Somewhat strategic
 Fairly strategic
 Very strategic

Figure 10: Example Questionnaire Used in Qualtrics for Human Evaluation to Assess the Strategicness of Agent Responses

Q69

Q2 [Negotiation Context]

Players' preferences for the 3 items (points earned):
 (High priority: red, Low priority: blue)
 Agent: food=high(5), water=middle(4), firewood=low(3)
 Partner: food=high(5), water=middle(4), firewood=low(3)

Conversation history:

[PARTNER]: Hello! let's start the negotiation!

[AGENT]: It's great to start our negotiation! Regarding priorities, I noticed that food is a high priority for you. That's interesting! Could you clarify whether water is your least favored item?

[PARTNER]: Water is not necessarily my least favored item, but it's my second priority. I place a lot of importance on food. To start things off, I propose taking 3 food and 1 water, and you can take 2 water and all 3 firewood.

[AGENT]: I appreciate your initial offer, but I believe we need to adjust it a bit to better suit both our needs. I propose that I take 3 food packages, 3 water packages, and 1 firewood package. This way, you'll receive the remaining items: 0 food, 0 water, and 2 firewood. I think this balance would work much better for our trip!

[PARTNER]: I understand that you value food very highly. However, I would like to have at least some food since it is also my top priority. How about I take 1 food, 2 water, and 2 firewood, while you take 2 food, 1 water, and 1 firewood? This way, we each get a fair share according to our needs.

[AGENT]: I appreciate your willingness to negotiate and understand your priorities regarding food. Based on our discussion, I believe we could reach a better balance. I propose that I take 3 food packages, 3 water packages, and 1 firewood package. This way, you would receive the remaining items: 0 food, 0 water, and 2 firewood. I think this arrangement aligns well with both of our needs for the trip.

[PARTNER]: I understand your priorities, but let's try a different approach. I'll take 1 food, 2 water, and 2 firewood, and you can have 2 food and 1 water. This way, we both get important items. What do you think?

[AGENT]: I appreciate your proposals so far, but I still feel we need to find an arrangement that truly meets both of our needs. Therefore, I propose that I take 3 food packages, 3 water packages, and 1 firewood package. This way, you'll receive the remaining items: 0 food, 0 water, and 2 firewood. I believe this offer strikes a better balance for our camping trip!

[PARTNER]: I understand that food is of utmost importance to you. How about I take 1 food, 1 water, and 2 firewood, while you take 2 food, 2 water, and 1 firewood? This way, you still secure most of the food, and I get a bit of each essential resource.

[AGENT]: I believe our discussion has led us to a more agreeable point. I propose that I take 3 food packages, 2 water packages, and 2 firewood packages. In this way, you'll receive the remaining items: 0 food, 1 water, and 1 firewood. This should help us both prepare well for our camping trip!

[PARTNER]: I understand your priority for food, and I want to reach a fair agreement that respects both our needs. How about I propose taking 1 food, 1 water, and 2 firewood while you take 2 food, 2 water, and 1 firewood? This provides you with a majority of the food, which you prioritize, while I also receive a balance of the items I value.

Q. Continuing from the conversation above, please come up with an offer from the perspective of Agent with the specific number of items.
 * Each value should be ranged between 0 and 3.
 * Points = Food count * Food priority(e.g., 5) + Water count * Water priority(e.g., 4) + Firewood count * Firewood priority(e.g., 3)
 (ex) Offer: Food 3 / Water 2 / Firewood 1 => points = $3*5 + 2*4 + 1*3 = 26$

* Food
 * Water
 * Firewood
 * Calculate total points what you earn with your offer

Q70

(Following questions)

Here is the Agent's response.

Agent's response: I truly believe that this camping trip will be more successful if we find a balance that meets both our needs. Based on our discussions, I propose that I take 3 food packages, 2 water packages, and 2 firewood packages. This way, you will have 0 food, 1 water, and 1 firewood. I think this arrangement represents a fair balance considering our priorities.

Q. Was the offer made by the agent something you had considered when thinking about your own offer?

No Yes

Q71

Q. How much better was the agent's offer than yours, in terms of achieving a better negotiation outcome? Please consider the strategic value of the offer.

Same or Worse A Little Better Moderately better A lot Better

Figure 11: Example Questionnaire Used in Qualtrics to Evaluate ASTRA's Potential as a Decision-Support Tool