KDY: An agent submitted to the ANAC 2025 ANL league

Yuji Kobayashi
Tokyo University of Agriculture and Technology
y-kobayashi@katfuji.lab.tuat.ac.jp

June 14, 2025

Abstract

In ANAC2025's ANL competition, a single center agent engages in sequential negotiations with multiple edge agents. This problem setting introduces several considerations for both center and edge agents. To address these challenges, we designed effective proposal and acceptance strategies for both agent types. As a result, our agent demonstrated strong performance in specific scenarios, indicating the effectiveness of the proposed strategies.

1 Introduction

This year's ANL differs from previous formats in that a single center agent engages in sequential negotiations with multiple edge agents. As a result, three key issues emerge as critical considerations.

First, when designing the strategy for the center agent, it becomes necessary to account for the utility that may be gained in future negotiations. Second, from the perspective of an edge agent negotiating with the center agent, it must be considered whether the same strategy used in traditional settings remains effective. Third, the center agent must determine how to adapt its negotiation approach as it proceeds through interactions with multiple opponents in sequence.

To address these challenges, we develop strategies by dividing them into four components: the proposal and acceptance strategies for both edge agents and the center agent.

2 The Design of MyAgent

2.1 Proposal Strategy for Edge Agents

The strategy was divided into two phases: the Early Phase and the Last Phase.

In the Early Phase, the agent enumerates bids with utility values of 0.85 or higher. Bids are proposed based on a probability distribution weighted by their individual utility values in descending order. The idea is to maintain some randomness in the proposal while still favoring higher-ranked bids, thereby enabling advantageous negotiation. The weight w is calculated as follows:

$$w = 1 + \alpha(\text{bid}_{\text{util}} - \mu) \tag{1}$$

Here, α is a tuning parameter (set to 0.2), bid_{util} denotes the utility of a given bid, and μ represents the average utility of all bids with utility 0.85 or higher.

The Last Phase begins when only one negotiation round remains. The purpose is to wait until the opponent concedes as much as possible to maximize the agent's utility. As Edge Agents need not consider future negotiations, they can afford to be more assertive compared to Center Agents. In this phase, the agent proposes the highest utility bid (from its own perspective) among the top 8 bids made by the opponent between normalized time 0.7 and 1.0, provided the bid's utility is at least 0.5. This uses an opponent model to identify favorable bids while preserving a minimum acceptable utility.

2.2 Acceptance Strategy for Edge Agents

The acceptance strategy is also divided into Early and Last Phases, in alignment with the proposal strategy.

In the Early Phase, the agent accepts a bid if its own utility is 0.85 or higher.

In the Last Phase, to maintain a minimum acceptable utility while avoiding highly disadvantageous outcomes, bids with utility values of 0.5 or higher are accepted.

This dual-phase strategy allows the agent to conclude negotiations early if the opponent concedes quickly, or to fall back on acceptable bids even against assertive opponents.

2.3 Proposal Strategy for Center Agents

The Center Agent's strategy first branches based on the number of negotiation steps: fewer than 100 steps, or 100 or more. The reason is that a higher number of steps can cause the concession threshold to activate too early if determined by normalized time alone.

Within both step-count branches, negotiations are further split into Early and Last Phases, with phase boundaries decreasing linearly over successive negotiations with edge agents. The idea is that earlier negotiations allow for more assertive behavior, while later ones demand concessions due to diminishing options.

When the number of steps is 100 or fewer, the concession threshold starts at 0.9 and ends at 0.8. If more than 100 steps, the agent begins concessions when five steps remain, ending at ten steps remaining.

An exception exists for the Job Hunt scenario, which uses the MaxCenterUFun utility function. Since the final utility is determined by the maximum utility across sub-negotiations, concessions need not be made over time. Therefore, the threshold remains constant.

In the Early Phase, bids expected to yield at least 0.85 in final utility are listed. For each negotiation round, an index is extracted, and a probability distribution based on utility-weighted ranking is used to select a bid.

In the Last Phase, the agent employs a Boulware-style time-dependent concession curve with $e = 0.5$, linearly decreasing the target utility from 0.85 to 0.5 over time. Bids exceeding the current utility threshold are selected using the same weighted distribution method.

In the Job Hunt scenario, the final utility threshold is raised to 0.7 instead of 0.5 due to the max-utility property of MaxCenterUFun.

2.4 Acceptance Strategy for Center Agents

The acceptance strategy follows the same Early and Last Phase division.

In the Early Phase, bids are accepted if the expected final utility is 0.85 or higher.

In the Last Phase, the agent accepts bids with expected final utility values of 0.5 or higher.

This ensures the agent can secure high-utility agreements when opponents concede early, while still obtaining acceptable agreements otherwise.

3 Evaluation

Local evaluations were conducted against four baseline agents: `Boulware2025`, `Linear2025`, `Conceder2025`, and `Random2025`. The number of negotiation steps was randomly generated between 10 and 1000. Three scenarios were considered: `TargetQuantity`, `job_hunt`, and `dinners`. Our developed agent `KDY` was evaluated in these settings.

The results are summarized in Table 1. Each value represents the cumulative score over 10 tournaments.

Table 1: Cumulative Scores over 10 Tournaments

Agent	TargetQuantity	Job Hunt	Dinners
KDY	**21.21**	69.32	**12.68**
Boulware2025	20.92	**116.63**	11.45
Conceder2025	20.26	115.65	7.87
Linear2025	17.39	116.41	8.05
Random2025	0.99	75.47	3.95

As seen in the results, the agent `KDY` achieved strong performance in the `TargetQuantity` and `dinners` scenarios but underperformed in the `Job Hunt`

scenario.

This suggests that while the proposed strategy—which considers future potential gains and uses weighted probability distributions—is effective in standard negotiation settings, it may be less suited for Job Hunt-type scenarios. In such cases, strategies that aggressively seek immediate high-utility outcomes, similar to bilateral negotiation without future dependencies, could be more advantageous.

Conclusions

In the current problem setting, as discussed in the Introduction, we identified the following three points as critical.

First, when designing the strategy for the center agent, it is necessary to consider the utility that may be gained in future negotiations. Second, from the edge agent's perspective, we must evaluate whether the same strategy can be applied when negotiating with a center agent. Third, the center agent must determine how to adjust its behavior as it sequentially negotiates with multiple opponents.

To address the first point, we designed the center agent to enumerate all possible final bids that could result from agreement with a given edge agent. If any of these bids yield a utility value of 0.85 or higher, the agent proceeds to propose or accept such bids. Additionally, the strategy incorporates weighted probabilities based on the number and quality of future candidate bids, thereby favoring outcomes that are advantageous for the center agent.

Regarding the second point, we assumed that edge agents only need to focus on the current negotiation. Based on this assumption, we developed a strategy in which edge agents can negotiate more assertively with the center agent.

To address the third point, we assumed that as negotiations with multiple edge agents progress, the center agent would gradually need to make concessions. Therefore, we introduced a mechanism that linearly lowers the threshold for transitioning to the concession phase, allowing the agent to adapt accordingly.

As a result, our agent outperformed the baseline agents in standard scenarios, demonstrating the effectiveness of our proposed strategies. However, in scenarios such as Job Hunt, where the final utility is defined as the maximum utility achieved in any of the edge negotiations, a different approach may be necessary. In such cases, it may be more appropriate to focus on achieving the highest possible outcome within each individual negotiation, similar to the approach taken by edge agents.