

OzuNegotiator: Deception-Aware Bilateral Negotiation for ANL 2026

Umair Ahmad¹ Vahit Pinar Eren²

^{1,2}Department of Computer Science, Özyğin University

¹umair.ahmad@ozu.edu.tr ²eren.pinar@ozu.edu.tr

Abstract

We present **OzuNegotiator**, a bilateral negotiation agent designed for the ANL 2026 league of ANAC 2026. The agent combines an ultra-stubborn two-segment Boulware aspiration curve with a four-phase deceptive bidding strategy and a frequency-based opponent model. Three novel mechanisms form the core of our deception strategy: a decoy issue trick that plants false signals in the opponent’s frequency model, bounded deceptive randomization that prevents preference inference from bid patterns, and opponent-aware top- k bid selection that steers late bids toward the Pareto frontier. In a full tournament across 7 scenarios and 504 negotiations against all ANL 2026 baseline agents, OzuNegotiator ranked **1st** with an ANL score of 1.472, outperforming BOANeg (1.388), MAP-Neg (1.388), SimpleNegotiator (1.179), BoulwareTBNegotiator (0.711), and ConcederTBNegotiator (0.467).

1 Introduction

In ANL 2026, two agents negotiate bilaterally using the Alternating Offers Protocol (AOP). Each agent knows only its own utility function; the opponent’s preferences must be inferred from observed bids. The ANL 2026 scoring rewards two components simultaneously:

$$\text{Score}_A = \text{adv}_A + \frac{\tau(A)}{\tau(A) + \tau(B)}$$

where $\text{adv}_A = (u_A(o) - r_A) / (\max u_A - r_A)$ is the normalized advantage, and $\tau(A) = \frac{1}{2}(1 + d(\hat{u}_B, u_B))$ is derived from the Kendall rank correlation between the opponent’s estimate of our utility function and our true one. Maximizing the score therefore requires both *high utility agreements* and *misleading the opponent’s model*.

2 Agent Architecture

OzuNegotiator follows the BOA (Bidding, Opponent modeling, Acceptance) architecture [1] with three novel

extensions targeting the deception component of ANL 2026.

2.1 Preprocessing

Before negotiation begins, all rational outcomes (those above the reservation value r) are pre-computed and sorted by descending utility:

$$R = \{o \in O \mid u(o) > r\}, \quad u(o_1) \geq u(o_2) \geq \dots \geq u(o_n)$$

This sorted list enables efficient $O(1)$ lookup at each turn without searching the full outcome space.

2.2 Two-Segment Aspiration Curve

The aspiration function $A(t)$ determines the minimum acceptable utility at time $t \in [0, 1]$:

Segment 1 ($t < 0.90$) — Ultra-stubborn Boulware:

$$A(t) = \max(r, r + (1 - r)(1 - t^{20}))$$

With exponent 20, the agent demands near-maximum utility for the first 90% of the negotiation ($t=0.8 \Rightarrow A \approx 0.988$), conceding only minimally.

Segment 2 ($t \geq 0.90$) — Linear deadline drop:

$$A(t) = \max\left(r, r + (1 - r)\left[p\left(\frac{1 - (t - 0.90)}{0.10}\right)\right]\right)$$

where $p = 1 - 0.90^{20} \approx 0.878$ is the pivot value. All meaningful concession is compressed into the final 10 rounds, preventing early exploitation while ensuring eventual agreement.

2.3 Four-Phase Bidding Strategy

At each turn, the agent forms the candidate set $C(t) = \{o \in R \mid u(o) \geq A(t)\}$ and selects from it based on the current phase:

Phase 1 ($t < 0.40$): Apply decoy filter to full $C(t)$, then sample uniformly at random. Maximizes entropy of our offer stream to prevent preference inference.

Phase 2 ($0.40 \leq t < 0.75$): Apply decoy filter to top 40% of $C(t)$, then sample uniformly. Narrows utility floor while maintaining deception.

Phase 3a ($0.75 \leq t < 0.90$): Select top 5 outcomes, then:

$$o^* = \arg \max_{o \in C_3(t)} \hat{u}_{\text{opp}}(o)$$

Phase 3b ($t \geq 0.90$): Select top 3 outcomes, same arg max selection. Protects against deadline exploitation — all late offers are near our maximum utility.

2.4 Decoy Issue Mechanism

The agent measures the importance of each issue as the utility range when that issue alone is varied:

$$\text{Importance}_i = U_{\max}(i) - U_{\min}(i)$$

The *least important* issue is designated as the decoy. Throughout Phases 1 and 2, only bids where the decoy issue equals a fixed target value are offered. Frequency-based opponent models [2] interpret consistent values as high importance, causing them to misrank our issue weights — reducing the Kendall rank correlation between their estimate and our true utility function, directly boosting our $\tau^{(A)}$ score.

2.5 Opponent Modeling

A frequency model tracks $\text{freq}[\text{issue}][\text{value}]$ for each opponent offer. Issue weights are computed as:

$$w_i = \frac{\text{max count for issue } i}{\text{total offers}}$$

and value utilities as $\hat{u}_i(v) = \text{freq}[i][v] / \sum_v \text{freq}[i][v]$. The estimated opponent utility is:

$$\hat{u}_{\text{opp}}(o) = \sum_i \frac{w_i}{\sum_j w_j} \cdot \hat{u}_i(o_i)$$

The model requires a minimum of 3 offers before use; prior to that, bid selection falls back to uniform random sampling. The model is used in Phases 3a/3b to select which of our top-utility bids the opponent is most likely to accept.

2.6 Acceptance Strategy

The agent accepts when any of the following holds:

1. $u(\text{offer}) \leq r \Rightarrow$ always reject (never irrational)
2. $u(\text{offer}) \geq 0.9 \cdot \max(u) \Rightarrow$ accept (near-perfect deal)
3. $u(\text{offer}) \geq A(t) \Rightarrow$ accept (meets aspiration)
4. $t > 0.97 \Rightarrow$ accept any rational offer (deadline safety net)

2.7 Deadline Middle-Offer Protection

At $t > 0.97$, if the gap between our last bid utility and the opponent’s offer utility exceeds 0.2, the agent proposes one compromise bid before accepting:

$$\text{target} = \frac{u(\text{last_bid}) + u(\text{opponent_offer})}{2}$$

choosing the rational outcome closest to this target from above. This prevents stubborn opponents from extracting a cheap deadline agreement.

3 Experimental Results

We ran a full tournament using an12026 tournament across 7 scenarios (Amsterdam, Camera, Car, Grocery, IS_BT_Acquisition, Laptop, NiceOrDie) against all 5 provided baseline agents, with each pair negotiating in both directions (504 total negotiations).

Table 1: Tournament results: average ANL 2026 score

Agent	Score	Rank
OzuNegotiator	1.472	1st
MAPNeg	1.388	2nd
BOANeg	1.388	3rd
SimpleNegotiator	1.179	4th
BoulwareTBNegotiator	0.711	5th
ConcederTBNegotiator	0.467	6th

OzuNegotiator ranked 1st with a score of 1.472. The 0.084 point gap over MAPNeg and BOANeg — both of which use opponent models — is attributable to the Concealing component: our decoy trick and randomized bidding actively distort the opponent’s frequency model, while MAPNeg and BOANeg make no effort to mislead their opponent. BoulwareTBNegotiator and ConcederTBNegotiator have no opponent model ($\tau = 0$) and therefore earn no deception points, explaining their substantially lower scores.

Per-scenario analysis shows OzuNegotiator achieves average Advantage above 0.88 in every tested domain (IS_BT_Acquisition: 0.935, Grocery: 0.930, Amsterdam: 0.925, Camera: 0.888, Car: 0.882), with 100% agreement rate across all matchups.

4 Conclusion

OzuNegotiator demonstrates that combining an ultra-stubborn aspiration curve with active deception mechanisms produces superior ANL 2026 scores compared to

agents that model the opponent but make no effort to mislead it. The decoy issue trick, bounded deceptive randomization, and opponent-aware late bidding are complementary: deception raises the Concealing score while opponent-awareness raises the Advantage score. Future extensions include explicit opponent behavior classification using the DANS framework [3] and extending the decoy mechanism to fool Bayesian opponent models.

References

- [1] T. Baarslag, K. V. Hindriks, M. J. C. Hendriks, A. S. Y. Dirkzwager, and C. M. Jonker. Decoupling negotiating agents to explore the space of negotiation strategies. In *Proc. ACAN 2012*, 2014.
- [2] O. Tunali, R. Aydogan, and V. Sanchez-Anguix. Rethinking frequency opponent modeling in automated negotiation. In *PRIMA 2017*, pp. 263–279. Springer, 2017.
- [3] K. Hindriks, C. M. Jonker, and D. Tykhonov. Let’s DANS! An analytic framework of negotiation dynamics and strategies. *Web Intelligence and Agent Systems*, 9:319–335, 2011.
- [4] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3):159–182, 1998.