

AOA007: A Deadline-Aware Adaptive Agent for Automated Negotiations

Achiya Zigler, Omer Silberman, Amit Katz

June 20, 2026

Abstract

The Automated Negotiation League (ANL) 2026 challenge requires agents to balance the pursuit of maximum utility with concealing their preferences from the opponent. In response, we present AOA007, a highly calculative agent built upon the MAPNegotiator architecture. Our agent utilizes a Monotonic Concession Rational Offering (MiCRO) strategy paired with a composite, time-aware acceptance threshold, an ensemble of six adaptive opponent models, and an iso-utility randomization layer for preference concealment. This multi-layered approach allows the agent to remain mathematically rigorous while actively confusing opponent modeling attempts, achieving strong performance on both the Advantage and Concealing dimensions of the ANL 2026 scoring metric.

1 Introduction

The Automated Negotiating Agent Competition (ANAC) fosters research into autonomous conflict resolution by presenting complex strategic environments. The ANL 2026 competition introduces a dual scoring metric: $\text{Score} = \text{Advantage} + \text{Concealing}$, where Advantage measures the utility gained above the reserved value, and Concealing measures how well the agent hides its true preferences from the opponent’s model (evaluated via Kendall’s τ rank correlation).

Our agent addresses both dimensions through a layered architecture. By leveraging the modularity of the NegMAS framework, we developed an agent that (1) operates with strict rationality to secure high-utility outcomes, (2) employs an emergency override mechanism to guarantee agreements when time is critically low, (3) uses an ensemble of time-adaptive opponent models to understand the opponent, and (4) applies iso-utility randomization to conceal its own preferences.

2 The Design of AOA007

The agent is constructed using the `MAPNegotiator` class, which serves as the coordinating framework for our specialized components. A critical tactical parameter is `acceptance_first=True`, which forces the evaluation of incoming offers before generating counter-offers. This ensures the agent never accidentally rejects a favorable deal by generating a counter-offer first.

2.1 Bidding Strategy

Our agent’s proposal strategy is governed by the `MiCROOfferingPolicy` (Monotonic Concession Rational Offering).

- **The Mechanism:** Instead of conceding along a smooth time-based curve, this policy maps the outcome space into distinct utility tiers. It systematically searches for the best possible outcome and only steps down to the next highest tier when mathematically forced to concede.
- **Strategic Advantage:** This calculative approach ensures the agent never yields “free” utility, maximizing profit potential against easily intimidated opponents during the first 98% of the negotiation.

2.2 Acceptance Strategy

The acceptance logic is the core of our adaptive approach, utilizing an `AnyAcceptancePolicy` to act as a logical OR-gate between two distinct rules:

- **Rule A: The Rational Baseline (ACNext)**

For the vast majority of the negotiation, the agent evaluates the opponent’s offer against its own planned trajectory. It accepts an offer only if the utility matches or exceeds the utility of the offer the `MicroOfferingPolicy` was about to propose. This ensures absolute logical consistency between our demands and our acceptances.

- **Rule B: The Deadline Panic (ACTime)**

To prevent walking away empty-handed, we integrated a strict time-based override. Normalized time (τ) is tracked continuously. This policy outputs a strict rejection until $\tau = 0.98$. Once 98% of the negotiation time has elapsed, the policy flips to “Accept” for any incoming offer. Because of our OR-gate architecture, this overrides the rational baseline and immediately triggers an agreement, securing a deal before the negotiation crashes.

2.3 Opponent Modeling

The agent employs an ensemble of six opponent models, four of which are custom time-adaptive extensions of standard algorithms. All models observe the opponent’s offer sequence and attempt to reconstruct their utility function. The ANL 2026 evaluator reads these models to compute the Concealing score.

- **ExponentialDecayFrequencyModel:** Extends `GSmithFrequencyModel` by weighting each observed offer with e^{5t} , where t is the relative time. Offers made near the deadline—when the opponent is under pressure to reveal true preferences—receive exponentially more weight than early exploratory offers.
- **TimeScaledHardHeadedModel:** Extends `GHardHeadedFrequencyModel` with a dynamic learning coefficient that starts at 0.05 and grows quadratically to 0.9 over time. This ignores early stubbornness (likely a decoy) and heavily rewards stubbornness near the deadline.
- **AdaptiveScalableBayesianModel:** Extends `GScalableBayesianModel` with a time-adaptive learning rate. Conservative (0.05) during the first 80% of negotiation, then jumps to 0.5 for the final phase when opponent offers become more revealing.
- **RuthlessAgentLGModel:** Extends `GAgentLGModel` by actively punishing issue weights (40% penalty) when the opponent changes a value between consecutive offers. Issues that remain stable are inferred to be high-priority.

- **GSmithFrequencyModel** and **GNashFrequencyModel**: Standard baselines providing frequency-based utility estimation without time adaptation.

The ensemble approach provides robustness: if one model is misled by a deceptive opponent, others may still capture the true preference structure.

3 Concealing Bidding Strategy

The ANL 2026 scoring metric rewards agents for hiding their preferences. A frequency-based opponent model observing a standard MiCRO agent would quickly identify preferred issue values from the deterministic offer sequence. To counter this, we introduce an **iso-utility randomization** layer.

3.1 Mechanism

During initialization, the agent pre-computes *iso-utility buckets*: groups of outcomes that provide approximately equal self-utility (within $\epsilon = 0.05$). When MiCRO selects an offer at utility level u , instead of proposing that specific outcome, the agent randomly samples from the bucket of all outcomes with utility in $[u - \epsilon/2, u + \epsilon/2]$.

3.2 Properties

- **Zero expected utility loss**: Since all outcomes in a bucket provide approximately the same self-utility, the agent sacrifices at most $\epsilon/2$ in any single offer.
- **Frequency flattening**: By randomizing among iso-utility alternatives, the agent’s offer history exhibits a near-uniform distribution over issue values, making frequency-based models unable to distinguish preferred from non-preferred values.
- **Bucket merging**: For utility levels with fewer than 5 outcomes, adjacent buckets are merged to guarantee meaningful randomization diversity.

3.3 Selection of ϵ

We tested $\epsilon \in \{0.005, 0.01, 0.02, 0.05, 0.10\}$ in tournament settings. The value $\epsilon = 0.05$ provided the best total score, as it creates large enough buckets for effective randomization while keeping utility loss negligible.

4 Evaluation

The agent was tested in local tournaments against 10 diverse opponents including standard benchmark agents (Boulware, Conceder, Linear, Aspiration), MiCRO-based agents, and a ClassicMAP agent equipped with frequency-based opponent models.

4.1 Quantitative Results

In a 12-competitor tournament across 7 scenarios (2016 total negotiations):

Agent	Advantage	Concealing	Score
AOA007 (ours)	0.682	0.709	1.391
ClassicMAP	0.658	0.703	1.361
MiCRONegotiator	—	—	0.759
BoulwareTBNeg.	—	—	0.659

Our agent consistently ranked first, outperforming both standard and model-equipped opponents.

4.2 Analysis

The agent excels in standard encounters. By remaining stubborn but rational, it frequently secures highly favorable agreements against conceding agents well before the panic threshold is reached. The iso-utility randomization provides a consistent concealing advantage, particularly against opponents equipped with frequency-based models.

A notable weakness emerges against identical hardliners. The agent routinely pushes these negotiations to the absolute limit. In these edge cases, the 98% panic trigger forces our agent to accept whatever sub-optimal deal is currently on the table, which occasionally results in lower final utility compared to earlier potential compromises.

5 Lessons and Suggestions

The development process highlighted several key dynamics in component-based automated negotiation:

1. **The Danger of Stubbornness:** Strict rationality is profitable but brittle. Without the ACTime safety net, our MiCRO strategy frequently timed out against tough opponents. Adaptation is strictly necessary for tournament survival.
2. **Randomness Beats Sophistication for Concealing:** We tested several concealing approaches—anti-frequency self-monitoring, decoy injection on low-weight issues, and temporal noise (non-monotone concession). All were outperformed by simple random sampling from iso-utility buckets. Deterministic concealing strategies create their own detectable patterns.
3. **Ensemble Opponent Models:** Time-adaptive models that discount early offers substantially outperform static frequency models. Opponents tend to explore broadly early and reveal true preferences under deadline pressure.
4. **Future Improvements:** Future iterations could (a) dynamically adjust the 98% panic threshold based on the opponent’s concession curve, (b) use the opponent model to select among iso-utility alternatives that are likely more acceptable to the opponent (model-aware offering), and (c) employ neural network-based opponent modeling trained on synthetic negotiation data.

Conclusions

AOA007 demonstrates a robust, modular approach to the ANL 2026 challenge. By combining the rigid rationality of MiCROofferingPolicy, the adaptive safety net of a composite acceptance strategy, an ensemble of time-adaptive opponent models, and iso-utility randomization for preference

concealment, the agent effectively addresses both dimensions of the ANL 2026 scoring metric. The architecture provides a strong foundation for managing time-pressured automated negotiations while actively resisting opponent modeling attempts.