

AgentNexus: A Privacy-Preserving Negotiation Agent for the ANL 2026 League

Nishanth Rajan

*Department of Geodesy and Geoinformatics, Leibniz Universität Hannover
Master of Science – Geodesy & Geoinformatics*

Abstract

The ANL 2026 competition challenges participants to design bilateral negotiation agents that simultaneously maximize agreement utility and conceal their own preferences from opponents. This dual objective creates a fundamental tension: high-utility offers cluster near the agent’s ideal point, which is exactly where an opponent’s model learns to look. This paper presents AgentNexus, a negotiation agent built on the NegMAS BOANegotiator architecture. AgentNexus addresses both objectives through four complementary mechanisms: (i) a three-phase Boulware concession schedule, (ii) a distribution-based frequency opponent model with dynamic issue-weight refinement and behavioral profiling, (iii) an active preference-masking strategy combining decoy signalling and masked window selection, and (iv) a multi-layer acceptance strategy with extreme-conflict detection and hardliner-advantage guards. Across 28 official ANL 2026 tournament runs between 8 and 20 June 2026, AgentNexus 2.0 placed first in 57.1% of runs and within the top three in 89.3% of runs, with a mean score of 5882.19 across all runs.

1 Introduction

Automated negotiation is a key area of multi-agent systems research. Agents are traditionally evaluated on the quality of agreements they reach. The ANL 2026 competition extends this by introducing a second, equally weighted objective: preference concealment. Agents are scored on both the utility of the final agreement and on how successfully they prevent the opponent from reconstructing their utility function.

Formally, the ANL 2026 score is defined as:

$$Score = (u(\omega) - r) + C$$

where $u(\omega)$ is the agent’s utility for the agreed outcome ω , r is its reservation value, and C is the Kendall- τ -based concealment score reflecting how much of the agent’s preference ranking the opponent’s frequency model failed to capture. Maximising both terms simultaneously is non-trivial because offers that maximise utility also concentrate near the agent’s ideal point, making that region easy for the opponent to identify.

AgentNexus manages this tension by separating offer generation into two stages: a utility-safe candidate set and a masking-optimal selection within that set. A three-phase Boulware concession schedule controls aspiration over time, while opponent behavioral profiling and a hardliner-advantage guard handle adversarial edge cases.

2 Agent Architecture

AgentNexus extends the BOANegotiator base class from NegMAS and overrides the single `__call__` entry point, giving full control over both acceptance decisions and offer generation. All tunable constants are collected in a frozen `NegotiationConfig` dataclass of approximately 60 parameters, making the agent reproducible and easy to tune. The four main components are:

- Distribution-based opponent model (`DistributionAnalyzer`).
- Three-phase Boulware concession and offer selection engine.
- Active preference-masking pipeline.
- Multi-layer acceptance strategy with conflict detection.

At the start of each negotiation, `on_preferences_changed` enumerates all outcomes above the reservation value, sorts them by own utility, and caches them in `_outcomes`. The estimated opponent utility function is published in `private_info` as required by the ANL 2026 scoring infrastructure.

3 Opponent Model

3.1 Distribution-Based Frequency Model

The DistributionAnalyzer maintains per-issue weighted frequency counts of opponent offers. For each issue i and value v , a smoothed value score is computed as:

$$\hat{u}_i(v) = (C_i(v) + \lambda) / (C^*_i + \lambda) ^{0.35}$$

where $C_i(v)$ is the occurrence count of value v on issue i , $\lambda = 1$ is a Laplacian smoothing factor, C^*_i is the maximum count over all values of issue i , and the power exponent 0.35 compresses high-frequency signals to reduce overconfidence.

The estimated opponent utility for an outcome ω is then the weighted sum of per-issue scores:

$$\hat{u}_{opp}(\omega) = \sum_i w_i \cdot \hat{u}_i(\omega_i)$$

Model update weights are not uniform. Earlier offers receive higher weight because they are closer to the opponent's genuine preference. Repeated identical offers are down-weighted exponentially with a damping factor of 0.48 and a minimum weight floor of 0.18, reducing the influence of stubborn repetition.

3.2 Dynamic Issue-Weight Refinement

Every $k = 4$ opponent offers, the agent compares the previous and current window distributions for each issue using L1 distance. If an issue's distribution remains stable across a window in which the opponent conceded on at least one other issue, that stability is interpreted as the opponent protecting a high-priority issue. Its weight is boosted by:

$$\Delta w_i = 0.08 \cdot (1 - t)$$

where t is the current relative time. The weight vector is renormalised after each update. Larger updates are therefore applied early in the negotiation when evidence is more informative.

3.3 Behavioral Profiling

Beyond frequency modeling, AgentNexus maintains four behavioral profile scores for the opponent throughout the negotiation:

- Hardliner score: high when the opponent shows low concession rate and consistently low offers for the agent.
- Conceder score: high when opponent offers are improving over time, estimated via pairwise concession slopes.
- Threshold-acceptor score: high when the opponent shows low offer diversity and stable self-utility values.
- Mirror-strategic score: high when the opponent shows diverse offers and intermediate utility levels.

These scores modulate concession aggressiveness, extreme-conflict rescue guards, and the emergency floor utility applied near the deadline.

3.4 Opponent Deadline Forecasting

To evaluate candidate offers under uncertainty, AgentNexus predicts the best utility the opponent may offer before the deadline. The forecast is computed by taking the median of all pairwise concession slopes from the offer history, scaling the resulting trend by the remaining time (capped at 0.35), and taking the maximum of the forecast and the best offer seen so far. A confidence score is derived from three factors: sample size (number of offers), time span covered, and offer diversity. The forecast and confidence together drive a one-step lookahead tiebreak score added to each candidate's composite score.

4 Concession Schedule

The target utility $T(t)$ at relative time $t \in [0, 1]$ follows a three-phase Boulware schedule with a floor of $\max(r, 0.64 u_{\max})$:

$$t < 0.55 \Rightarrow T = 0.92 \cdot u_{\max} \quad (\text{early, aggressive})$$

$$0.55 \leq t < 0.85 \Rightarrow T = 0.92 u_{\max} - 0.17 u_{\max} \cdot ((t-0.55)/0.30)^{2.4} \quad (\text{mid, Boulware})$$

$$t \geq 0.85 \quad \Rightarrow \quad T = 0.75 \, u_{\max} - 0.11 \, u_{\max} \cdot ((t-0.85)/0.15)^{1.4} \quad (\text{late, final})$$

In the early phase the agent maintains a high aspiration of $0.92 \, u_{\max}$. The mid phase applies a Boulware-shaped concession with exponent 2.4, meaning the agent holds close to its aspiration for most of the phase and concedes only near the transition. The late phase applies a final concession with a softer exponent of 1.4 down to a minimum of $0.75 \, u_{\max}$ before the floor applies.

An opponent-responsive adjustment is overlaid on this schedule from $t = 0.45$ onward. When sufficient model confidence exists, the target is nudged upward if the opponent is conceding (rewarding patience) or downward if the opponent is a hardliner with urgency building (bounded hold of 5.5% $u_{\max}$ and concession of 4.5% $u_{\max}$ respectively). On the final step, the target receives an additional relief of 2.5% $u_{\max}$ to facilitate last-second agreements.

5 Preference-Concealing Bidding Strategy

5.1 Issue-Cost Estimation

At initialisation, AgentNexus estimates each issue's cost as the range of own utility across its values, normalised by $u_{\max}$. Issues with low cost are “cheap” (own utility varies little across their values) and issues with high cost are “expensive” (own utility is sensitive to their value). This classification drives the decoy signalling mechanism.

5.2 Decoy Signalling

During the early and mid phases ($t < 0.78$), each candidate offer is scored by a decoy signal:

$$s_{\text{decoy}} = 0.65 \cdot \text{cheap_stability} + 0.35 \cdot \text{expensive_variety}$$

`cheap_stability` rewards keeping cheap-issue values unchanged from the previous offer, creating the appearance of resistance on dimensions that cost little. `expensive_variety` rewards changing expensive-issue values between offers, hiding the dimensions on which the agent's utility is most sensitive. Together these signals mislead the opponent's frequency model about which issues matter most.

5.3 Masked Window Selection

After multi-objective scoring, all candidates whose own utility falls within a slack band around the best candidate are collected into a masked window. The slack is 3.5% of $u_{\max}$ in the early phase and 7.0% in the late phase. Within this window, candidates are re-ranked by a masking score:

$$s_{\text{mask}} = 0.55 \cdot \text{issue_novelty} + 0.30 \cdot (1 - \text{value_reuse}) + 0.15 \cdot (1 - \text{exact_repeat})$$

`issue_novelty` measures the fraction of issue values that differ from the previous offer. `value_reuse` measures how frequently the proposed values appeared in recent offers. `exact_repeat` penalises offers identical to a recent bid. The window is capped between 5 and 17 candidates. Final selection within the window uses a deterministic step-based index to avoid pseudo-random patterns that opponents could exploit.

5.4 Two-Level Candidate Filtering

From the mid game ($t \geq 0.60$), the candidate pool undergoes a two-level filter before masking. First, only outcomes within 3.5% $u_{\max}$ of the best available utility are retained, preserving own utility. Second, these safe candidates are ranked by estimated acceptance likelihood and the top 60% are kept. This ensures that the agent only presents offers that are both utility-safe and likely to be accepted, without sacrificing concealment.

5.5 Multi-Objective Candidate Scoring

Each candidate offer is scored as:

$$s = w_{\text{own}} \cdot s_{\text{own}} + w_{\text{acc}} \cdot s_{\text{acc}} + w_{\text{par}} \cdot s_{\text{par}} + w_{\text{mask}} \cdot s_{\text{mask}} + 0.04 \cdot s_{\text{deadline}} + 0.035 \cdot s_{\text{decoy}} + 0.025 \cdot s_{\text{lookahead}}$$

Phase-specific weight vectors (w_{own} , w_{acc} , w_{par} , w_{mask}) are: early (0.74, 0.06, 0.04, 0.16), mid (0.60, 0.16, 0.10, 0.14), late (0.48, 0.28, 0.16, 0.08). The Pareto score s_{par} balances a Nash-like product, social welfare, and offer balance. The deadline score s_{deadline} scales opponent utility by urgency beyond $t = 0.85$.

6 Acceptance Strategy

Acceptance is evaluated in strict priority order:

- Reservation guard. Any offer below the reservation value r is immediately rejected.
- Extreme conflict rescue. If the domain is very constrained (fewer than 5% of outcomes above $0.70 u_{\max}$) and the negotiation is stalled, a rescue compromise is computed. The opponent's best offer is accepted if it is within 2% of this compromise. This rescue is bypassed entirely when the hardliner advantage guard is active, preventing exploitation by stubborn opponents.
- ACcombi. At $t \geq 0.99$, the offer is accepted if its utility meets or exceeds the window-average of recent opponent offers.
- Target utility. Accept if $u(\text{offer}) \geq T(t)$.
- Last-chance. On the final negotiation step, accept if the offer is at least as good as the best opponent offer seen so far and meets $0.98 T(t)$.
- Counter-offer margin. Accept if $u(\text{offer}) \geq (0.995 - 0.08t) \cdot u(\text{planned counter-offer})$.
- Late-game floor. At $t > 0.97$, accept any offer above the emergency floor of $0.55 u_{\max}$ (raised to $0.62 u_{\max}$ when the hardliner advantage guard is active).

7 Evaluation

AgentNexus 2.0 (submission ID 21639) was submitted repeatedly to the official ANL 2026 competition server between 8 June and 20 June 2026, across 28 completed tournament runs. Across all runs, the agent finished first in 16 tournaments (57.1%) and placed within the top three in 25 tournaments (89.3%), with a mean score of 5882.19. Four of these runs (tournaments 19032–19038) were evaluated against a larger 20,000-outcome domain pool and are not directly comparable in absolute score to the remaining 24 runs, which used a 10,000-outcome pool and produced a mean score of 5189.34 with a first-place rate of 54.2% and a top-three rate of 91.7%.

The single best score recorded was 11149.38 (tournament 19034, large pool), though this run ranked only 7th, indicating that absolute score is highly dependent on the opponent pool and scenario set for that run rather than being directly comparable across tournaments. Within the more frequent 10,000-outcome pool, the agent's best result was 6416.40 (tournament 19115, rank 5) and its most consistent strong performances clustered around 5000–5600 with first-place finishes, for example tournament 19114 (5607.37, rank 1) and tournament 19164 (5383.21, rank 1).

Table 1: Full official ANL 2026 tournament history for AgentNexus 2.0 (ID 21639), 8–20 June 2026.

Tournament	Date	Score	Rank	Pool size
19190	2026-06-20	5092.32	2	10,000
19175	2026-06-20	5185.57	2	10,000
19164	2026-06-19	5383.21	1	10,000
19154	2026-06-19	5113.32	3	10,000
19144	2026-06-19	4710.85	1	10,000
19142	2026-06-18	5007.07	1	10,000
19141	2026-06-19	5007.07	1	10,000
19139	2026-06-18	5189.31	1	10,000
19136	2026-06-18	4868.87	1	10,000
19124	2026-06-18	5406.58	1	10,000
19115	2026-06-15	6416.40	5	10,000
19114	2026-06-14	5607.37	1	10,000
19111	2026-06-13	5158.90	2	10,000

Tournament	Date	Score	Rank	Pool size
19097	2026-06-12	4935.28	2	10,000
19087	2026-06-11	5020.79	1	10,000
19083	2026-06-11	4340.86	1	10,000
19080	2026-06-10	5329.75	1	10,000
19077	2026-06-10	4536.60	3	10,000
19074	2026-06-10	4649.74	2	10,000
19055	2026-06-10	5003.72	2	10,000
19047	2026-06-09	5088.60	1	10,000
19043	2026-06-09	4700.05	2	10,000
19041	2026-06-09	5383.34	1	10,000
19038	2026-06-08	10562.87	1	20,000
19036	2026-06-08	9702.54	1	20,000
19035	2026-06-08	8742.31	1	20,000
19034	2026-06-08	11149.38	7	20,000
19032	2026-06-08	7408.69	4	10,000

Rank variance across runs reflects differences in the opponent pool sampled for each tournament instance rather than agent instability: the same agent version (0.87.44) was used throughout this period without modification. The overall first-place and top-three rates suggest that AgentNexus performs robustly across a wide range of opponent compositions, though individual tournament rank should be interpreted alongside pool size and opponent diversity rather than score alone.

8 Discussion

Three design decisions contributed most to the agent’s performance.

First, separating utility-safe candidate generation from masking-optimal selection allowed the agent to pursue concealment without sacrificing agreement quality. The masked window mechanism ensures that offers within a narrow utility band are selected for their deceptive properties rather than pure utility rank, directly improving the Concealing component of the score.

Second, the hardliner advantage guard prevented the extreme-conflict rescue from being exploited. In early versions without this guard, the rescue mechanism conceded too deeply to opponents who never intended to move, resulting in agreements well below the Pareto frontier. Gating the rescue on a composite hardliner-mirror-conceder profile eliminated these losses.

Third, requiring minimum model confidence before activating opponent-responsive concession avoided erratic mid-game behaviour caused by noisy early evidence. The median-based deadline forecast was more robust to outlier offers than mean-based alternatives.

Potential limitations include additional computational overhead from per-candidate multi-objective scoring in large outcome spaces, and sensitivity of the behavioral profiles to opponents with mixed strategies that do not fit the four archetypes cleanly. Future work may investigate learned opponent type classifiers and multi-round preference evidence aggregation across repeated interactions.

9 Conclusion

This paper presented AgentNexus, a privacy-preserving negotiation agent for the ANL 2026 competition. The agent combines a distribution-based frequency opponent model with dynamic issue-weight refinement, a

three-phase Boulware concession schedule, an active preference-masking pipeline of decoy signalling and masked window selection, and a multi-layer acceptance strategy with hardliner-aware conflict resolution. Across 28 official ANL 2026 tournament runs, AgentNexus 2.0 placed first in 57.1% of runs and within the top three in 89.3%, with a mean score of 5882.19. The results demonstrate that explicit preference-concealment mechanisms, when integrated into a principled bidding architecture, can provide a meaningful and consistent competitive advantage in environments where opponent inference plays a central role.

References

- [1] T. Baarslag, K. Hindriks, C. Jonker, S. Kraus, and R. Lin. The first automated negotiating agents competition (ANAC 2010). In *New Trends in Agent-Based Complex Automated Negotiations*, Springer, 2012.
- [2] O. Tunali, T. Aydogan, and M. Sanchez-Anguix. Rethinking frequency opponent modeling in automated negotiation. In *Proceedings of PRIMA 2017, LNCS vol. 10621*, Springer, 2017.
- [3] T. Baarslag, K. Hindriks, M. Hendriks, A. Dirkzwager, and C. Jonker. Decoupling negotiating agents to explore the space of negotiation strategies. In *Proceedings of the 5th International Workshop on Agent-based Complex Automated Negotiations (ACAN)*, 2012.
- [4] Y. Farouq et al. NegMAS: A platform for automated negotiation, supply chain management and simulation. In *Proceedings of the 12th International Workshop on Agent-Based Complex Automated Negotiations (ACAN)*, 2019.
- [5] R. Aydogan, T. Baarslag, K. Hindriks, M. Jonker, and C. M. Jonker. A baseline for non-linear bilateral negotiations: The full results of the agents competing in ANAC 2014. In *Modern Approaches to Agent-based Complex Automated Negotiation*, Springer, 2017.
- [6] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
- [7] J. F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.