

ChangAgent: A Privacy-Aware Adaptive Negotiator for ANL 2026

Agent module: `chang_agent` | Agent class: `ChangAgent`

Abstract. ChangAgent is a bilateral multi-issue negotiation agent designed for the ANL 2026 privacy challenge. The agent seeks high agreement utility while making its revealed bid sequence less informative about its true utility function. It uses a bounded-regret acceptance policy, an adaptive bathtub-shaped aspiration curve, and a three-expert opponent model combining Smith-style frequency estimation, distribution-based frequency updates, and recency/transition evidence. The design emphasizes robust performance against heterogeneous opponents rather than exploiting a single benchmark family.

1. Competition Setting and Design Objective

ANL 2026 evaluates autonomous agents in bilateral multi-issue negotiations under the Alternating Offers Protocol. Each agent knows its own utility function and reservation value, but not the opponent utility function. If no agreement is reached, each side receives its reservation value. The score combines normalized agreement advantage with a privacy-related bonus based on how difficult it is for the opponent to reconstruct the agent's preferences. Therefore, an effective strategy must solve two partially conflicting tasks: reach agreements with high self utility, and avoid producing a bid sequence that directly exposes the true preference ordering [1].

ChangAgent treats the task as a controlled information-release problem. It maintains a strong self-utility objective but changes how offers are selected within a utility band. Instead of always proposing the highest-ranked remaining outcome, it chooses among multiple strategically equivalent candidates using opponent-model support, novelty, exact opponent-offer anchors, and privacy/diversity terms. The agent stores no cross-negotiation information; all learning is confined to the current session, consistent with the rules of participation [1].

2. High-Level Architecture

Input signals	Internal modules	Decision outputs
Own utility function, reservation value, deadline, offer history, and opponent responses.	Opponent model ensemble, adaptive aspiration controller, candidate generator, privacy scorer, and bounded-regret acceptor.	Accept, reject and counteroffer, or final concession offer.

The implementation follows the standard Bidding-Opponent modeling-Acceptance decomposition commonly used in automated negotiation, while using a custom single-file NegMAS negotiator for submission compatibility. NegMAS provides the underlying Python negotiation platform and supports SAO-style negotiators, time-based agents, MiCRO agents, tit-for-tat agents, and Genius-style baselines, which makes it suitable for local robustness tests [3].

3. Concealing Bidding Strategy

ChangAgent uses an adaptive bathtub aspiration curve. The target utility starts high, drops to a negotiated valley in the middle phase, and partially rises again near the deadline if the opponent's behavior indicates that stronger offers may still be accepted. The middle valley is not fixed. It is inferred from the opponent's historical proposals using a weighted combination of recency, repeated occurrence, and opponent-model support. This prevents the agent from conceding to a static low target when the opponent has already shown willingness to offer better outcomes.

At each proposal step, the agent builds a candidate pool from several sources: self-best outcomes near the current aspiration level, outcomes previously proposed by the opponent, Pareto-like outcomes with high estimated opponent value, probing offers for early learning, and decoy offers that preserve self utility while varying exposed issue-value patterns. Candidate scoring uses self utility, robust opponent utility, distance from the current target, novelty, whether the exact outcome was proposed by the opponent, and privacy/diversity terms. The exact-score structure is implementation-dependent, but conceptually follows:

```
score(o) = self_term(o) + opponent_term(o) + target_closeness(o) + novelty(o) +
opponent_anchor(o) + privacy(o) - repetition_penalty(o)
```

The privacy term favors offers that increase the entropy of the observed issue-value sequence without sacrificing too much self utility. This is intended to make simple frequency-based inference less reliable: when several outcomes have similar self utility, the agent avoids repeatedly exposing the same top issue values. Repetition is not forbidden, because repeating an opponent-seen anchor can improve agreement probability; instead, repetition is controlled and contextual.

4. Opponent Modeling

The opponent model is an online ensemble that returns an estimated opponent utility together with a rough uncertainty and confidence value. This is necessary because the ANL 2026 concealing score depends on opponent utility estimation quality; an agent without a functioning opponent model receives no share of the corresponding extra point [1]. The ensemble has three experts:

Expert	Signal used	Purpose
Smith-style frequency model	Counts of values proposed by the opponent and stability of issues across offers.	Produces an early estimate from limited evidence. It assumes frequently proposed values are likely important.
Distribution-based frequency model	Changes in value distributions across small non-overlapping windows of opponent offers.	Raises issue weight for issues that stay stable while the opponent concedes elsewhere, inspired by Tunali, Aydogan, and Sanchez-Anguix (2017) [4].
Transition/recency model	Recent offers, value transitions, and persistence of issue choices.	Reduces overreaction to early anchors and gives more influence to recent credible behavior.

During bidding, ChangAgent usually uses a risk-adjusted estimate rather than the raw mean: the estimated opponent utility is discounted when the experts disagree. This avoids overfitting a proposal to a fragile model. The same opponent model is exposed through the agent's opponent-utility adapter so that the evaluation system can request an opponent utility estimate.

5. Bounded-Regret Acceptance Strategy

The acceptance policy was designed after observing that a naive late-stage relaxation can increase agreement rate while reducing advantage too much. ChangAgent therefore uses segmented acceptance with a regret bound. Early in the negotiation, it accepts only offers at least as good as the utility of its next planned proposal. In the middle phase, the acceptance threshold is the maximum of a reservation-safe floor, a target-related floor, and a best-seen anchor. This allows strong offers to be accepted but avoids accepting a much worse offer immediately after receiving a better one.

Near the deadline, the agent becomes more willing to accept offers below its next planned proposal, but the relaxation is bounded. The late threshold remains above reservation plus a safety margin and is anchored to the best offer received so far. A forecast shortcut can accept an offer slightly above threshold when the expected benefit of waiting is weak. This makes the final phase deadline-aware without turning it into unconditional acceptance.

Phase	Acceptance behavior	Risk controlled
Early	Accept only if the offer is no worse than the agent's own next proposal.	Prevents premature reveal and low-utility agreements.
Middle	Use target floor, reservation safety, and best-seen anchor.	Avoids giving away utility after the opponent has already offered better deals.
Late	Permit bounded regret relative to next proposal and best-seen anchor.	Balances agreement probability against final advantage loss.

6. Tuning and Evaluation Methodology

The package includes a local tuning script that evaluates a grid of parameter presets by repeatedly calling the ANL runner. The main tunable parameters are the valley height, late-phase start, acceptance floor, middle regret, late regret, and final regret. Tuning is not used during an official negotiation; it only selects static defaults before submission and avoids cross-negotiation learning.

To reduce overfitting, candidate parameters should be tested against heterogeneous opponents: time-based Boulware, Linear, and Conceder agents; tit-for-tat agents; MiCRO-style rational concession agents; random or nice/tough agents; and BOA/MAP examples from the ANL template. The recommended metric is not only mean score but also agreement rate, advantage conditional on agreement, and score variance across domains. A parameter set that wins only against a narrow set of local opponents is less reliable than one with slightly lower peak performance but stable results across opponent families.

7. Limitations and Future Work

ChangAgent intentionally avoids complex learning that would be brittle under short deadlines and unknown domains. Its opponent model is approximate and frequency-based; it does not reconstruct a full probabilistic utility distribution. Its privacy control is also heuristic: it diversifies visible issue-value patterns but cannot guarantee low Kendall correlation against all adversarial estimators. Future work should consider a controlled portfolio of acceptance modes, calibration of opponent-model uncertainty from local traces, and automated parameter selection using nested validation over held-out domains and opponent families.

References

- [1] ANL 2026 Call for Participation. Automated Negotiation League 2026: An Overview. <https://scml.cs.brown.edu/files/anl/y2026/2026cfp.pdf>
- [2] ANL 2026 League Tutorial. <https://scml.cs.brown.edu/files/anl/y2026/template2026.pdf>
- [3] NegMAS Documentation. <https://negmas.readthedocs.io/en/latest/>
- [4] Tunali, O., Aydogan, R., and Sanchez-Anguix, V. (2017). Rethinking Frequency Opponent Modeling in Automated Negotiation. PRIMA 2017, LNCS 10621. DOI: 10.1007/978-3-319-69131-2_16.

Submission metadata: Agent Name = ChangAgent; Agent Module = chang_agent; Agent Class = ChangAgent.