

BetterCallAgentInfinityV1000: Concealment-Aware Time-Dependent Negotiation for ANL 2026

Steffan Wolter, Johnny Tarbouch, Qasim Sefaldeen
Leibniz University Hannover

Abstract. BetterCallAgentInfinityV1000 is a rule-based ANL 2026 negotiator for bilateral multi-issue alternating-offers negotiation under a concealing-preferences scoring rule. The agent treats the competition objective as a multi-objective control problem: maximize normalized advantage above reservation value, maintain agreement probability, and reduce the information leaked by its own bidding trace. The design combines a Boulware-style aspiration function, a bounded Nash-inspired concession floor, a time- and evidence-weighted opponent-utility proxy, and a Kendall Tau-b based concealment penalty. This report formalizes the implemented policy, maps its components to the submitted code, and identifies the design assumptions that determine its expected behavior.

1. Competition setting and objective


ANL 2026 evaluates agents in bilateral multi-issue negotiation using the Alternating Offers Protocol (AOP): one party proposes an outcome and the other accepts, rejects with a counteroffer, or walks away before the deadline [1]. Each agent observes its own utility function and reservation value, but not the opponent's. The distinctive 2026 challenge is concealment: an agent is rewarded not only for the agreement it reaches, but also for how well its preferences remain hidden from the opponent [1].

Let Ω be the outcome space, $u_i: \Omega \rightarrow \mathbb{R}$ agent i 's utility, r_i its reservation value, and o the final agreement. ANL normalizes the agreement component as advantage $adv_i = (u_i(o) - r_i) / (\max u_i - r_i)$. Concealment is evaluated through an opponent-model estimate: the Kendall Tau-b rank correlation $d(u_{\hat{j}}, u_j)$ is normalized as $\tau_j = 0.5(1+d)$. The additional concealment point is divided proportionally between the agents' normalized distances, with $\tau=0$ if no opponent model is provided [1]. Thus the agent must optimize agreement utility while controlling how much its own offers reveal.

$$score_i = adv_i + \tau_i / (\tau_i + \tau_j), \quad \tau_i = 0.5 * (1 + d(u_{\hat{j}}, u_j))$$

2. Design rationale and relation to prior work

The design follows the decomposition used in automated negotiation research: a bidding strategy, an opponent model, and an acceptance condition, as emphasized by BOA-style analyses of negotiation agents [5,6]. Time-dependent concession tactics, including Boulware behavior, originate from the negotiation decision functions of Faratin, Sierra, and Jennings [4]. Opponent modeling is treated conservatively because learned opponent utilities are known to be uncertain in bilateral negotiation [7]. The concealment penalty uses Kendall's rank correlation coefficient as a direct proxy for the same rank-based idea used by the ANL 2026 scoring rule [1,9].



Per-round control loop implemented in BetterCallAgentInfinityV1000

Figure 1: Architecture of the per-round decision loop. The agent filters rational outcomes, constructs a bounded representative subset, scores candidate counteroffers with own utility, opponent-utility proxy, and leakage penalty, then applies acceptance or counteroffer logic. The diagram reflects the submitted implementation rather than a generic negotiation architecture.

3. Formal policy

At initialization the agent evaluates the available outcomes and constructs the rational set $R = \{o \text{ in } \Omega \mid u(o) > r\}$. Outcomes are sorted by own utility. For large domains, the agent builds a representative subset S of size at most $K5=240$ by combining high-utility rational outcomes with randomized coverage of the remaining outcome space. This subset is used for repeated Pareto/Nash-inspired calculations and for estimating the concealment proxy at bounded cost.

The aspiration curve is a Boulware-style threshold. Let $u_{\max}$ be the best rational utility and F the current concession floor. The implemented aspiration is:

$$A(t) = u_{\max} - (u_{\max} - F) * t^e, \quad F = \max(r, \alpha * N_{\text{hat}})$$

Here t in $[0,1]$ is relative time, e is a domain-adaptive exponent, α is a floor multiplier, and N_{hat} is a Nash-inspired target computed over S using own utility and an online proxy for opponent utility. Since the true opponent utility is unobserved, the target functions as a bounded bargaining heuristic rather than an exact Nash bargaining solution.




Figure 2: Time-dependent aspiration and late recovery. The dark-blue curve is the Boulware aspiration $A(t)$, the LUH-blue curve is the acceptance threshold, and the amber dashed line is the bounded floor F . The red shaded region marks the very late safety-net window, where strictly positive-advantage offers are accepted to avoid preventable no-agreement outcomes.

The opponent-utility proxy combines an inverse-utility prior with a HardHeaded frequency model. Before evidence accumulates, the prior assumes partial preference opposition, $1-u(o)$. As opponent offers arrive, the learned frequency score $HH(o)$ receives increasing weight:

$$u_{\text{opp_hat}}(o, t) = (1-\lambda_{\text{t}})(1-u(o)) + \lambda_{\text{t}} HH(o), \quad 0 \leq \lambda_{\text{t}} \leq 0.92$$

λ_{t} increases with both the number of observed opponent offers and relative time. The model is used as a ranking signal, not as a hard truth, which limits overfitting to noisy early offers.

4. Concealment-aware bidding and acceptance

The core contribution is the offer objective. For each candidate o , the agent evaluates own utility $u(o)$, estimated opponent utility $u_{\text{opp_hat}}(o)$, a balance term $\min(u(o), u_{\text{opp_hat}}(o))$, and a leakage proxy $L(o)$. The leakage proxy is based on the trace of issue values already offered by the agent. Once enough own offers exist, the agent computes a Kendall Tau-b relation between candidate-induced frequency scores and true own utility over the representative subset S . A high alignment would make the agent's private ranking easier to infer, so it is penalized.

$$q(o) = h(t)L(o) - s(t)(u(o)+u_{\text{opp_hat}}(o)) - 0.08 \min(u(o), u_{\text{opp_hat}}(o)); \text{ choose } \operatorname{argmin} q(o)$$

The time weights $h(t)$ and $s(t)$ create a negotiation trade-off. Early in the session the concealment penalty has more influence; later, agreement pressure rises and utility terms dominate. This implements the ANL objective more directly than a pure Boulware agent, which might achieve high utility but reveal its preferences by repeatedly offering the same top-ranked issue values.




Figure 3: Objective-level view. The left panel summarizes the minimized counteroffer score $q(o)$: own utility, opponent proxy, and balance lower the objective, while rank-leakage raises it. The right panel maps the agent objective to the official ANL 2026 scoring structure: normalized advantage plus a concealment share derived from Kendall Tau-b opponent-model quality.

The acceptance condition combines reservation safety, aspiration-based acceptance, and late risk control. Offers below reservation are never accepted. Otherwise, the agent accepts if the offer clears the current aspiration-adjusted threshold. Very near the deadline, it accepts any strictly positive-advantage offer and may re-propose the best acceptable opponent offer seen so far. This rule is intentionally delayed to reduce exploitability while preventing no-deal outcomes that are worse than a modest agreement.

5. Implementation mapping

Component	Implementation	Role in the report
Rational set	_a3, _a4, _a5	Utility-sorted outcomes above reservation value.
Representative subset	_m0, K5=240	Bounds runtime for Nash-inspired floor and leakage estimation.
Aspiration	_m3, _m4, _m5, _m8	Boulware curve with bounded Nash-inspired concession floor.
Opponent proxy	_C1, GHardHeadedFrequencyModel, _m1, _m2	Mixture of inverse-utility prior and frequency model.
Leakage penalty	_b2, _m10, kendalltau	Penalizes rank-revealing offer-frequency traces.
Acceptance/recovery	_m6, _m7, _m11, K9=0.985	Combines threshold acceptance with late positive-deal recovery.

Table 1: Mapping from report concepts to the submitted implementation. The table is intended to make the scientific description auditable against the code.

6. Parameters, complexity, and reproducibility

Constant	Value	Interpretation
K0	0.73	Concession-shape weight used in aspiration and acceptance factors.
K1	0.988	Base time gate for late acceptance/reproposal.
K2	0.82	Opponent-model learning-rate factor.
K3	0.84	Base concealment weight in candidate scoring.
K4	36	Maximum candidates scored per round after filtering.
K5	240	Representative subset cap.
K6	7	Refresh period for concession floor recomputation.
K7	10	Small-domain threshold.
K8	10.0	Base Boulware exponent.
K9	0.985	Safety-net activation time.

Table 2: Principal policy constants. These are reported explicitly because they determine concession speed, model trust, runtime bounds, and late risk posture.

Operation	Approximate cost	Bounded by
Initialization	$O(\Omega \log \Omega)$ when enumerated	Outcome-space enumeration and sorting.
Candidate scoring	$O(K4 * K5)$ in leakage phase	Candidate cap K4 and subset cap K5.
Floor refresh	Library-dependent Pareto/Nash over S	Subset cap K5, refreshed every K6 steps.
Memory	$O(R + K5 + \text{issue-value counts})$	Rational set plus compact frequency trace.

Table 3: Runtime considerations. The agent is engineered for tournament use by bounding repeated expensive operations, while still preserving high-utility candidates.

The implementation uses Python 3.11.8 with NegMAS, NumPy, and SciPy. It uses no large language model, no external service, no persistent cross-negotiation memory, and no offline training data. Randomization is seeded from deterministic domain properties to introduce candidate diversity without relying on hidden state. A reproducible empirical evaluation consists of the official ANL/NegMAS tournament harness with fixed domains, opponent sets, seeds, and ablation variants, allowing utility, agreement rate, concealment score, and runtime to be compared under a single protocol.

7. Limitations

- The inverse-utility prior can fail in domains where preferences are positively correlated rather than opposed.
- Frequency-based opponent modeling can be manipulated by deceptive or non-stationary opponents.
- The Kendall Tau-b leakage proxy protects primarily against rank/frequency inference, not arbitrary opponent inference models.
- The representative subset can miss important Pareto/Nash candidates in very large outcome spaces.
- The late safety net improves agreement robustness but can be exploited by opponents that intentionally delay acceptable offers.

8. Conclusion

BetterCallAgentInfinityV1000 is best understood as a concealment-aware Boulware negotiator. Its main methodological contribution is not any single component, but the coupling of a time-dependent aspiration curve with a bounded Nash-inspired floor, a gradually trusted opponent proxy, and a Kendall Tau-b leakage penalty aligned with the ANL 2026 concealing objective. The approach is lightweight and auditable, while its limitations identify clear directions for ablation-based evaluation.

References

- [1] ANAC organizers. Automated Negotiation League (ANL) 2026: Call for Participation. <https://anac.cs.brown.edu/files/anl/y2026/2026cfp.pdf>
- [2] ANAC organizers. Automated Negotiation League (ANL). <https://anac.cs.brown.edu/anl>
- [3] T. Baarslag et al. The Automated Negotiating Agents Competition, 2010-2015. *AI Magazine* 36(4), 115-118, 2015. DOI: 10.1609/aimag.v36i4.2609.
- [4] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation Decision Functions for Autonomous Agents. *Robotics and Autonomous Systems* 24(3-4), 159-182, 1998. DOI: 10.1016/S0921-8890(98)00029-3.
- [5] T. Baarslag et al. Decoupling Negotiating Agents to Explore the Space of Negotiation Strategies. *ACAN* 2012.
- [6] T. Baarslag et al. The Significance of Bidding, Accepting and Opponent Modeling in Automated Negotiation. *Frontiers in AI and Applications*, 2014. DOI: 10.3233/978-1-61499-419-0-27.
- [7] T. Baarslag et al. Learning about the Opponent in Automated Bilateral Negotiation: A Comprehensive Survey of Opponent Modeling Techniques. *AAMAS* 30(5), 849-898, 2016. DOI: 10.1007/s10458-015-9309-1.
- [8] Y. Mohammad, S. Nakadai, and A. Greenwald. NegMAS: A Platform for Automated Negotiations. *PRIMA 2020, LNCS 12568*, Springer, 2021.
- [9] M. G. Kendall. A New Measure of Rank Correlation. *Biometrika* 30(1-2), 81-93, 1938. DOI: 10.1093/biomet/30.1-2.81.