

Badiron: A Firm-Concession, Pareto-Bidding Agent with First-Appearance Opponent Modelling for ANL 2026

Team 408 - Automated Negotiation League (ANAC 2026) - Agent: BadIronV4 (module badiron, class Badiron)

Abstract

Badiron is a fixed-heuristic bilateral negotiator for the ANL 2026 challenge, whose score rewards two things at once: individual ADVANTAGE (agreement utility above the reservation value) and a relative OPPONENT-MODELLING / CONCEALING contest scored by Kendall rank correlation. The agent couples (i) a firm time-based concession curve with an AC-next acceptance rule and a late endgame deal-securing sweep, (ii) a live opponent model that ranks each issue value by WHEN IT FIRST APPEARS in the opponent's offers - exploiting the fact that a conceding opponent reveals its preferred values first - and republishes it every turn as a linear-additive utility function, and (iii) Pareto-frontier bidding that, among the outcomes still acceptable to us, always offers the one the opponent is estimated to value most. Against a varied field of optimised NegMAS strategies on both zero-reserve and reservation-varying scenarios, Badiron improves mean normalised advantage by 5-8% over the previous version while holding the concealing contest at parity, and ranks first on the local competition-faithful benchmark.

1. Problem and Scoring

ANL 2026 is a bilateral negotiation run under the Stacked Alternating Offers Protocol (SAO) over a fixed, known number of rounds ($n_steps=100$). Each side has a private linear-additive utility function and a reservation value; failing to agree yields the reservation value. The per-negotiation score has two additive halves:

$$\begin{aligned} \text{score} &= \text{advantage} + \text{concealing_share} \\ \text{advantage} &= (u(\text{deal}) - \text{reserve}) / (\text{max_u} - \text{reserve}) \quad \text{in } [0,1] \\ \text{concealing_share} &= \tau_{\text{self}} / (\tau_{\text{self}} + \tau_{\text{opp}}) \\ \tau &= (1 + \text{kendall}(\text{true_ufun}, \text{model})) / 2 \end{aligned}$$

Here τ_{self} measures how well WE model the opponent and τ_{opp} how well they model us; the concealing half is therefore a relative contest split between the two negotiators. Three levers move the score: (1) reach high-utility agreements, (2) model the opponent accurately, and (3) make our own preferences hard to recover. A critical NegMAS contract is that the scorer reads our model from `private_info['opponent_ufun']`; if it is never set, τ_{self} is forced to zero and the entire concealing half is silently forfeited. Badiron republishes a real linear-additive utility function under that key on every turn.

2. Strategy

2.1 Concession and acceptance

The aspiration (target) utility follows a firm Boulware curve, $\text{target}(t) = \text{max} - (\text{max} - \text{floor}) * t^e$ with $e=12$, so the agent holds its demands near the maximum for most of the negotiation and concedes only late. The floor is adaptive: it never drops below a fraction (0.85) of the best offer the opponent has already made us, since that utility is

effectively on the table. Acceptance combines three rules: AC-next (accept any offer at least as good as the utility we would demand next step); immediate acceptance of near-optimal offers (normalised advantage ≥ 0.98), which captures excellent, possibly non-recurring offers; and an endgame regime - after 97% of the rounds we sweep the offer down toward the floor and, in the final steps, accept any positive-advantage deal rather than time out into a zero-advantage disagreement.

2.2 Opponent modelling: first-appearance ordering

Because the opponent concedes from its best outcomes toward its worst, the ORDER in which a value first appears in its offers is a strong signal of preference: the earlier a value first appears, the more the opponent likes it; values never offered are scored lowest. Badiron assigns each value a score of $1 - \text{first_appearance_step} / \text{offers_seen}$, adds a boost to the values in the opponent's very first offer (its revealed ideal point), and weights each issue by one minus the normalised entropy of its observed value distribution (issues the opponent holds firm on are important). These value scores and weights define a linear-additive utility function that is republished every turn. In an offline study over 154 recorded negotiations, first-appearance ordering reached Kendall optimality ~ 0.78 versus ~ 0.73 for dense-rank frequency; alternative schemes (raw frequency, recency weighting, Herfindahl / dominance / deviation-time issue weights, stickiness) were all weaker.

2.3 Pareto-frontier bidding

Each turn the agent takes the set of outcomes still acceptable to it (utility $\geq \text{target}$) and offers the one the opponent model values most, breaking ties deterministically toward the outcome best for us. Every concession therefore stays on the efficient frontier: the opponent accepts sooner, on terms that remain good for us. This directly buys advantage and - because the same model also drives modelling - improves both halves of the score at once. It also spreads our offers over many outcomes, a mild free concealment effect, instead of repeating a single bid.

3. Design Methodology and Key Findings

The agent was tuned with a fast head-to-head harness that pits it against a varied field of optimised NegMAS strategies (Boulware, Conceder, Linear, Aspiration, Hybrid, MiCRO, NaiveTitForTat, TopFraction) plus model-publishing agents (a BOA agent, a multi-model agent, and a naive 'reverse' model), recomputing the exact ANL 2026 metric per negotiation. Several findings shaped the final design:

- Test on reservation-varying scenarios. The bundled domains all have a zero reservation value, which is not representative; conclusions flip once reservations are non-zero. We evaluate on a seeded mix spanning reservation levels in addition to the real domains.
- Firmness wins the advantage half. A high Boulware exponent ($e=12$) beat softer curves: against opponents that concede, firmness extracts more, and against firm opponents advantage is low for everyone regardless.
- Only OUR advantage matters, not the opponent's. The score is our advantage plus our concealing share; the opponent's utility does not enter, so 'being exploited' is irrelevant and securing any positive deal beats a zero disagreement.
- A hard minimum-offer floor hurt. Refusing to concede below a fixed advantage threshold lost score; sweeping fully toward the floor late captures low-but-positive deals that beat disagreements.

- First-appearance modelling beats frequency, and a deterministic, advantage-optimal tie-break removes ~ 0.03 of run-to-run variance that pure randomisation introduced with no concealment benefit.

4. Results

On a representative set combining the real domains with reservation-varying generated scenarios, Badiron improves mean normalised advantage by roughly 5-8% over the previous version while holding the concealing share at parity (~ 0.80 locally, where few opponents publish models). On the competition-faithful local benchmark (the official anl2026 metric, rotated ufuncs, $n_steps=100$) it ranks first in its field with 100% agreement, no runtime errors, and high outcome efficiency (Pareto-optimality ~ 0.96). The behaviour-based contract suite (valid offers; 1-4 issues; a published, non-degenerate opponent model against every opponent and ufunc ordering) passes in full.

5. Limitations and Future Work

The local opponent field is necessarily weaker and less model-heavy than the live competition, so the measured gains are a conservative lower bound: the two upgrades (Pareto bidding and first-appearance modelling) specifically target the regimes a strong field stresses - pickier opponents that reward good bids, and many opponents that contest the concealing half. The clearest remaining headroom is a more powerful opponent model (e.g. Bayesian or rank-learning over the full offer sequence) to dominate rather than split the concealing contest, and opponent-type-adaptive concession. No information is persisted between negotiations (per the competition rules); all learned state is per-instance and initialised when preferences are set.

6. Conclusion

Badiron is a compact, robust heuristic that pushes all three score levers, advantage-first. Firm concession with a late deal-securing endgame, Pareto-frontier bidding driven by a first-appearance opponent model, and always publishing that model under the scored key together yield a measurable improvement over the prior version while remaining fast and error-free.

References

[1] R. Aydogan et al. The Automated Negotiating Agents Competition (ANAC). [2] NegMAS: a Python framework for automated negotiation, <https://negmas.readthedocs.io>. [3] T. Baarslag et al. Acceptance conditions in automated negotiation. [4] ANL 2026 Call for Participation (2026cfp).