

StaBornNeg: An agent submitted to the ANAC 2026 ANL league

Athina Georgara , Sarvapali D. Ramchurn
University of Southampton, Southampton, UK
{a.georgara, sdr1}@soton.ac.uk

1 Introduction

The 2026 ANL challenge asks participants to build a *bilateral, multi-issue* negotiation agent that, while operating under the Alternating Offers Protocol (AOP), **conceals its own preferences** from the opponent. An agent’s final score has two additive components:

1. the normalized **advantage** $\text{adv}_A = \frac{u_A(o) - r_A}{\max(u_A) - r_A}$, the utility of the agreed outcome o relative to the reservation value r_A ; and
2. a shared **concealing point**, split between the two agents according to how poorly each side’s opponent model estimates the other’s true utility function. The estimate quality is measured by the (normalized) Kendall τ -b rank correlation between the true and estimated utility functions.

The implication is that an agent must simultaneously (a) reach a good agreement and (b) behave in a way that makes its own preference *ranking* hard to infer. **StaBornNeg** (a “stubborn”, static-born negotiator) is a deliberately minimal, *constant-bid* agent built around a single anchor outcome. The three required behavioural components are described below.

2 Acceptance Strategy

The acceptance logic is implemented in `acceptance_strategy` and is driven by a single pre-computed reference outcome, `self._reveal_offer`. This anchor is selected once during initialization (`on_preferences_changed`) as the rational outcome whose utility is the largest value that does *not exceed* 85% of the agent’s maximum utility:

$$\text{_reveal_offer} = \arg \min_{o : u(o) \leq 0.85 u_{\max}} (0.85 u_{\max} - u(o)).$$

In effect, it is the agent’s “aspiration level” fixed at roughly 85% of the best attainable utility. The agent accepts the opponent’s current offer in either of two cases:

1. **Aspiration rule (any time):** accept if $u(\text{offer}) \geq u(\text{_reveal_offer})$, i.e. the incoming offer is at least as good as the 85% anchor. This is a flat, time-independent threshold.
2. **Deadline safety-net:** if the negotiation is almost over ($\text{relative_time} > 0.98$) and the offer merely clears the reservation value ($u(\text{offer}) > r$), accept it to avoid walking away with only r .

Otherwise the offer is rejected. The strategy is therefore a *constant aspiration* acceptor with a last-moment fallback: it holds out for $\geq 85\%$ utility for almost the entire horizon, then, in the final 2% of rounds, relaxes to accept anything better than no agreement. There is no gradual time-based concession between these two regimes.

3 Concealing Bidding Strategy

The bidding logic in `concealing_bidding_strategy` is intentionally trivial: every counter-offer the agent makes is the *same* outcome, `self._reveal_offer` (the 85%-of-max anchor). The opening offer, and every subsequent counter-offer, is this single bid.

```
def concealing_bidding_strategy(self, state):
    return self._reveal_offer # always the fixed 85%-of-max anchor
```

Why a constant bid conceals. The concealing score depends on how well the opponent can reconstruct the *ranking* of the agent’s outcomes from observed behaviour. Most opponent models (including a frequency/concession model) infer preferences from *how an agent concedes over time*: the sequence of bids traces a path down the agent’s utility ordering, revealing which issues and values it values most. By emitting a *single, unchanging* offer, `StaBornNeg` provides no concession trajectory at all. The opponent observes exactly one point of the agent’s utility space, repeated indefinitely, and therefore has almost no signal from which to recover the full ranking. This is expected to drive the opponent’s Kendall τ estimate toward chance, increasing `StaBornNeg`’s share of the concealing point.

Trade-off. The cost of this approach is bidding inflexibility. Because the agent never concedes and never tailors bids toward the opponent’s apparent preferences, agreement is only reached when the opponent itself concedes up to the 85% anchor (or, near the deadline, when the agent accepts via the safety-net). The design thus trades bidding sophistication and joint-gain search for strong, simple concealment.

4 Opponent-Modelling Strategy

Although **StaBornNeg** does not *exploit* its opponent model when bidding, it still builds one. The model is a classic **frequency (counting) model** that is updated every time an opponent offer arrives (`update_opponent_model`) and materialized as a `LinearAdditiveUtilityFunction` (`.build_opponent_ufun`).

Value frequencies. For each issue, the agent maintains a count of how often the opponent has proposed each possible value (`_value_counts`). The estimated value score for value v on an issue is its count normalised by the most frequent value’s count, so the opponent’s most-offered value on each issue receives a score of 1.0:

$$\hat{V}_{\text{issue}}(v) = \frac{\text{count}(v)}{\max_{v'} \text{count}(v')}.$$

Before any data arrives, every value is given a flat score of 1.0, so the model starts uninformative.

Issue weights via concentration. The relative importance (weight) of each issue is estimated from the **Herfindahl concentration** of its offer distribution,

$$H_{\text{issue}} = \sum_v \left(\frac{\text{count}(v)}{\text{total}} \right)^2,$$

which lies in $[1/k, 1]$ for k values. The intuition: if the opponent repeatedly insists on the *same* value for an issue (high concentration), that issue is probably important to them, so it receives a higher weight; if their offers spread evenly across values (low concentration), the issue is treated as less important. The raw weights H_{issue} are then normalized to sum to 1. The resulting linear-additive function is stored in `self.private_info["opponent_ufun"]`.

Summary

StaBornNeg is a deliberately stubborn agent. It (1) accepts on a fixed 85%-of-max aspiration threshold with a deadline safety-net above the reservation value; (2) conceals its preferences by repeating one constant anchor bid, denying the opponent any concession trajectory to learn from; and (3) models the opponent with a normalised frequency-count value function weighted by per-issue Herfindahl concentration, satisfying the requirement to produce an estimate and thereby remain eligible for the concealing point. The design optimises heavily for the concealment half of the ANL 2026 score at the expense of adaptive bidding, and its most natural avenues for improvement are a time-dependent concession schedule and using the opponent model to steer offers toward mutually agreeable outcomes without leaking the agent’s own ranking.