

SBDANL: An agent submitted to the ANAC 2026 ANL league

Your Team Member names, affiliations and emails

Abstract

We present SBDANL, our entry to the ANAC 2026 Automated Negotiation League, whose challenge is to maximise negotiation advantage while *concealing* one’s own preferences from the opponent. Our central empirical finding is that, against frequency-based opponent models, preference concealment and advantage are *intrinsically coupled*: utility is essentially a function of which issue-values one offers, so any high-advantage bid necessarily exposes one’s high-value options. Deliberate concealment tricks (oscillating the target utility, or flattening one’s offered-value histogram) therefore buy very little concealment for a large advantage cost. SBDANL instead takes the only *free* concealment available—selecting, on a late-conceding Boulware iso-utility band, the outcome the opponent likes most—and concentrates its effort on the two levers that actually move the score: a high but robust concession exponent and a frequency opponent model. It wins a round-robin against all competition sample agents and standard NegMAS baselines.

1 Introduction

In ANL 2026 two agents negotiate over a discrete multi-issue domain under the Alternating Offers Protocol for a fixed number of rounds. Each agent knows only its own utility function u and reservation value r . An agent’s score is

$$\text{score} = \underbrace{\frac{u(o) - r}{\max u - r}}_{\text{advantage}} + \underbrace{\frac{\tau_{\text{self}}}{\tau_{\text{self}} + \tau_{\text{opp}}}}_{\text{concealing share}}, \quad \tau = \frac{1}{2}(1 + \text{KendallTauB}(\hat{u}, u)), \quad (1)$$

where $\hat{u}$ is an agent’s estimate of the opponent’s utility function and the concealing share is split by how well each side models the other. Two consequences shaped our design. First, an agent that produces *no* opponent-model estimate scores $\tau_{\text{self}} = 0$ and forfeits the concealing point entirely; a working opponent model is mandatory. Second, the concealing share is at most 1, so it can be as valuable as the advantage term—if one can lower τ_{opp} without sacrificing advantage. We show this is largely impossible, and design accordingly.

2 The Design of SBDANL

SBDANL follows a Bidding–Opponent-model–Acceptance decomposition on top of SA0CallNegotiator. On preference initialisation it enumerates the rational outcomes (utility above r) and caches their normalised utilities. Each turn it (i) updates a frequency opponent model from the received offer, (ii) accepts if the offer clears a Boulware acceptance threshold, otherwise (iii) proposes a counter-offer chosen on an iso-utility band. The opponent-model estimate is published in `private.info["opponent_ufun"]` so that it contributes to τ_{self} . Crucially, the agent stores no state across negotiations, as required by the rules.




Figure 1: Pushing concealment effort up (histogram flattening, x -axis) fails to hide us better—the opponent’s model of us (dashed) does not fall below 0.5—and the total score (red) decreases. “Greedy iso-utility bidding” ($\text{conceal_weight} = 0$) is the best concealment that does not cost advantage.

3 Concealing Bidding Strategy

The offer target follows a Boulware aspiration that stays near the maximum early and concedes towards r only near the deadline: $\text{target}(t) = 1 - (1 - r_n)t^e$ in normalised utility, with concession exponent $e > 1$. Among rational outcomes whose utility lies in a thin band $[\text{target}, \text{target} + \beta]$ (protecting advantage), the agent proposes the outcome that its opponent model rates highest.

This last choice is our concealment mechanism, and it is essentially free: by selecting *within our own iso-utility set according to the opponent’s preferences*, the value pattern of our offers is driven by the opponent rather than by us, decorrelating our offer stream from our true ranking. Empirically this already pushes the opponent’s estimate of us to $\tau_{\text{opp}} \approx 0.49$ (i.e. no correlation).

We also implemented and tested the two “obvious” concealment ideas and report them as negative results (Section 6). Oscillating the target utility downward in large waves (the eponymous “snake”/*guwan-guwan* curve) does lower τ_{opp} slightly, but only by genuinely offering low-utility outcomes, which costs far more advantage than it gains. Actively flattening our own offered-value histogram is worse than useless: it *raises* τ_{opp} . The reason is fundamental—utility is approximately a monotone function of which issue-values one concedes, so a high-utility (good-advantage) bid is necessarily built from our preferred values and cannot hide them. Both knobs are retained in the code but disabled by default ($\text{wiggle} = 0$, $\text{conceal_weight} = 0$).

4 Acceptance Strategy

Acceptance is decoupled from the (potentially oscillating) offer target and measured against the smooth Boulware base $b(t) = 1 - (1 - r_n)t^e$: we accept an offer whose normalised utility is at least $b(t)$. A final-moment clause accepts any offer above the reservation value once $t > 0.97$. This end-game grab is what makes a high concession exponent safe: even against a never-conceding opponent the two clauses guarantee we take any rational agreement rather than walk away, preventing mutual-no-agreement collapse.

5 Opponent Model

We use a Smith-style frequency model. For each issue, a value’s score is its (recency-weighted) frequency in the opponent’s offers, normalised per issue; the issue weight is estimated from value *stability*—issues whose value rarely changes between consecutive offers are deemed important. Early offers are up-weighted ($w = 1 + \lambda(1 - t)$) since an opponent reveals its ideal first. The resulting `LinearAdditiveUtilityFunction` is rebuilt each turn. An A/B comparison (Section 6) shows stability weights plus recency beat a variance-based, no-recency model, but only marginally: τ_{self} saturates near 0.54 because opponents reveal few distinct offers—the bottleneck is information, not the estimator.

6 Evaluation

All experiments use the seven competition scenarios plus randomly generated ones, run as cartesian round-robins with the official `anl2026` opponent-modelling metric.

Round-robin ranking. Against every sample agent shipped with the skeleton and several Neg-MAS baselines, SBDANL ranks first (Table 1).

Agent	Mean ANL2026 score
SBDANL (ours)	1.497
BOANeg (sample)	1.411
MAPNeg (sample)	1.411
SimpleNegotiator (sample)	1.290
BoulwareTBNegotiator	0.611
HybridNegotiator	0.607
MiCRONegotiator	0.594
NaiveTitForTatNegotiator	0.469
MyNegotiator (skeleton)	0.399
ConcederTBNegotiator	0.373

Table 1: Round-robin over all built-in plus 12 generated scenarios.

Concession-exponent robustness. Sweeping e against a soft-to-tough panel (Conceder, Linear, Boulware, Tough, Tit-for-Tat, Hybrid, BOA, MAP) shows that higher e is robustly better and does *not* collapse against stubborn opponents (thanks to the end-game grab). The panel-mean score peaks around $e = 30$, but $e \geq 50$ starts losing to reciprocating opponents (Tit-for-Tat falls from 1.58 at $e=10$ to 1.29 at $e=50$) and degrades self-play agreement (from 100% to 91%). We therefore set $e = 20$, retaining about 99% of the peak with far better behaviour against stubborn/reciprocal agents.

Negative results on concealment. As Fig. 1 shows, neither the “snake” utility-wiggle nor histogram flattening improves the total score; both reduce it. The free concealment of greedy iso-utility bidding ($\tau_{\text{opp}} \approx 0.49$) is the practical optimum.

7 Lessons and Suggestions

The key lesson is the *advantage–concealment coupling*: because utility is largely determined by the issue-values one offers, a self-interested high-utility bidder cannot simultaneously hide its preferences from a frequency model. Concealment is therefore not a free dimension to optimise; the largest gains come from advantage (a robust, late concession) and from a working—if information-limited—opponent model. We suggest the organisers consider whether the concealment incentive is strong enough to overcome this coupling, and that future entrants treat “smart iso-utility selection” as the default rather than chasing explicit obfuscation. Promising future work for us is a better opponent model that exploits the *order* in which an opponent concedes, and scenario classification to set e per domain.

Conclusions

SBDANL is a deliberately simple, robust agent built on a single empirical insight: in ANL 2026 you cannot cheaply hide your preferences, so you should take the free concealment of greedy iso-utility bidding and otherwise maximise advantage with a robust Boulware concession and a frequency opponent model. It wins a round-robin against all sample agents and standard baselines.