Looking Ahead to Win: A Utility Fit Agent for Sequential Multi-Deal Negotiations

Garrett Seo   Tri-an Nguyen   Xintong Wang

Problem Setting

Sequential Multi-deal Negotiation

We consider a one-to-many negotiation setting where a center agent (one) negotiates with multiple edge agents (many) in a fixed sequence.

For each one-to-one negotiation between the center and edge agent, or a **subnegotiation**, the two agents follow a bilateral, **alternating-offers protocol**. Here, both agents take turns proposing offers and choosing to reject or reach an agreement.

The center agent is rewarded for the combination of all agreements reached from every subnegotiation.

A Scheduling Example

Consider the following scheduling example where Claire, the center agent, negotiates dinner with two friends, A and B (edge agents).

Figure 1. Source: ANL 2025 Call For Participation

The challenge here is that agreements in earlier subnegotiations can affect the final outcome and reward. So, at each subnegotation, **what is the best agreement and how can we realize that agreement?**

Lookahead Tree for Estimated Utilities

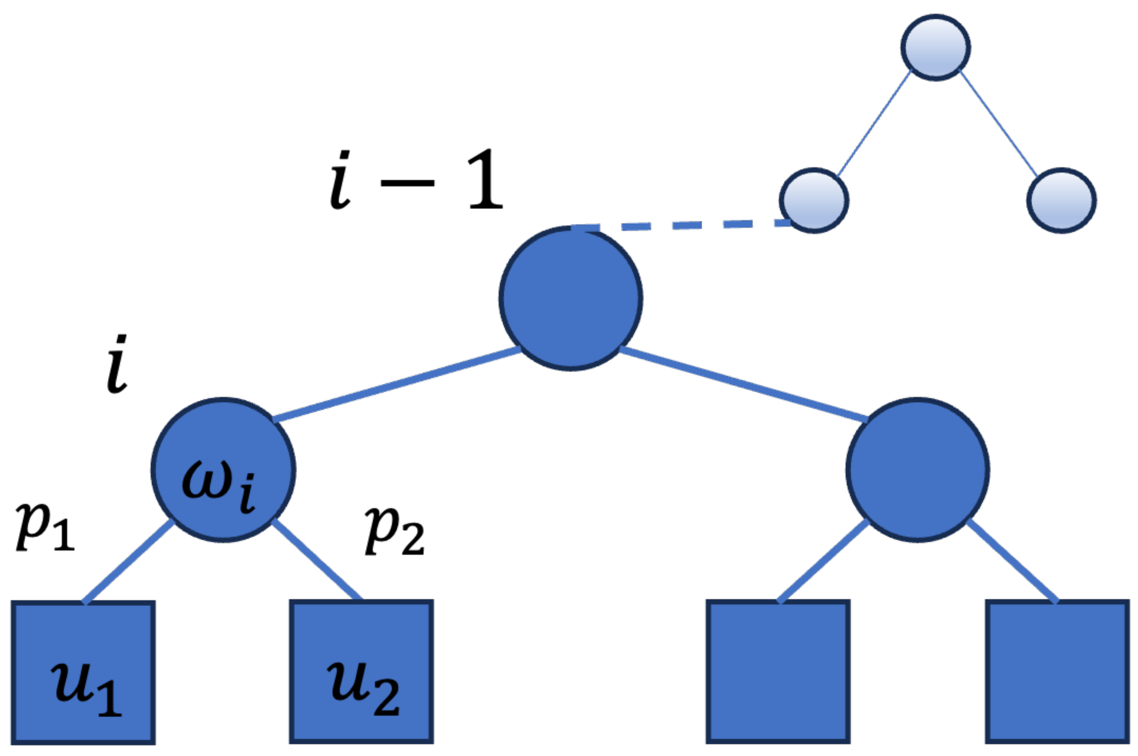

Figure 2. Recursive Estimation for Node ω_i. $\tilde{u}(\omega_i) = p_1 u_1 + p_2 u_2$. This value is propagated back to the node at depth $i-1$.

Tree Representation

- A node at depth i represents the the beginning of subnegotiation i. It contains all suboutcomes from previous subnegotiations $0, 1, 2, ..., i-1$.
- Children represent all suboutcomes of subnegotiation i.
- **Goal**: Solve for estimated utilities of suboutcomes using tree at each subnegotiation.

Recursive Estimation

1. Calculate estimated utility of all children.
2. Assign probability to children.
3. Propagate expected utility of parent upwards.

Modeling Opponent's Strategy

Family of Time-Based Strategies (Faratin et al. 1998)

$$x(t) = u_{min} + (u_{max} - u_{min}) * (1 - t^{\beta})$$

- x(t): **opponent's** utility of bid offered at time t.
- $u_{min/max}$: opponent's minimum/maximum utility achievable.
- β: opponent's concession degree.

Key Assumption: Opponents will follow this time-based strategy.

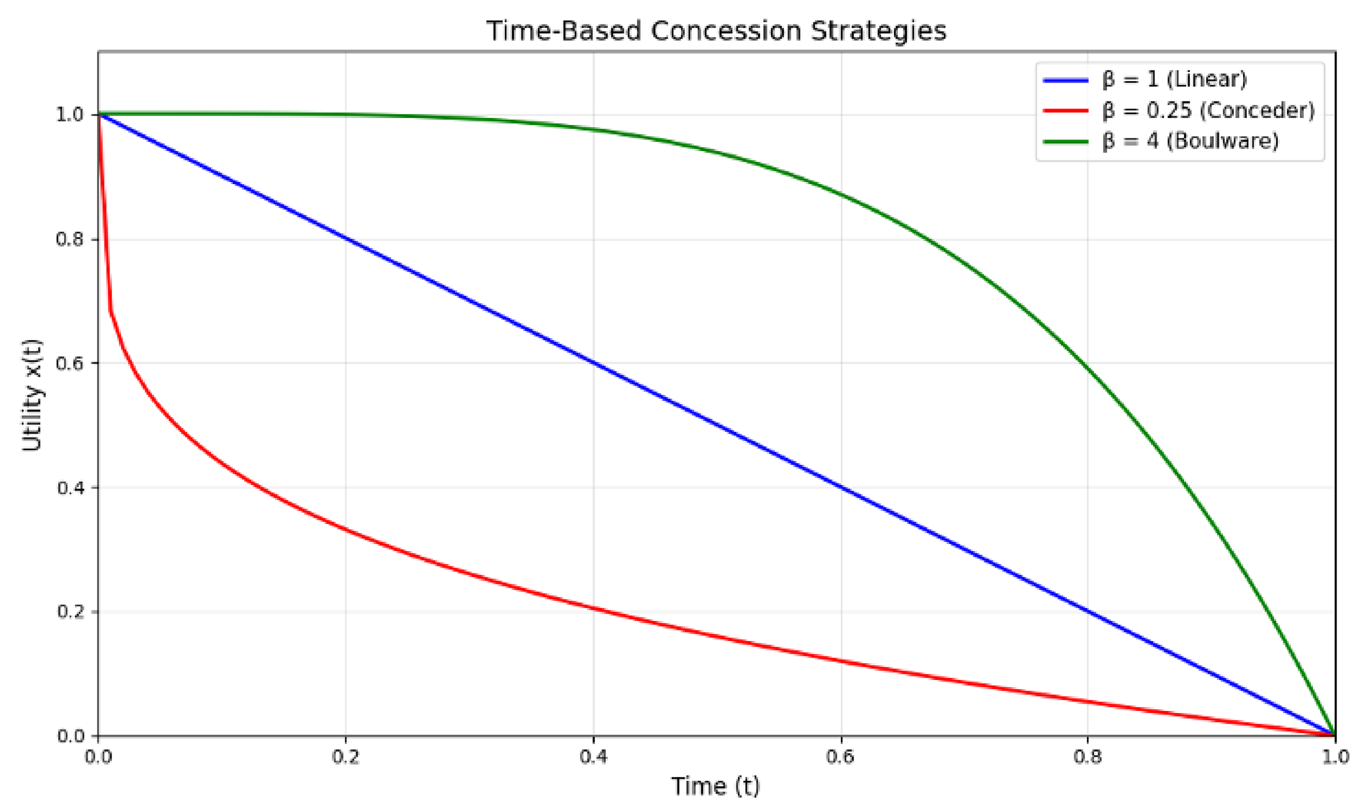

Our Concession Strategy

Due to no discount, we delay the negotiation to gain information from opponent's offers to make the most informed bid.
Key Assumption: Opponents are adversarial. Thus, as opponents concede over time, our utility on opponent offers should increase.
Utility Fit: We map opponent offers to the estimated utilities from the lookahead. Use these offers to fit a utility curve to reach a "projected estimated utility" that we bid at the final timestep.

Utility Fit

Goal: Fit parameters $\mathbf{u_{max}}, \boldsymbol{\beta}$ to the following function:

$$x'(t; u_{max}, \beta)$$

- x': **our estimated utility** from opponent's offers.
- u_{max}: our maximum utility the opponent is willing to concede to.
- β: opponent's concession degree.

Final timestep: Propose outcome with estimated utility from lookahead, $x'(t=1)$