

An Advanced Hybrid Strategy of *Oneshot* Agent for SCML OneShot: Cautious Partner Selection and Scalable Offer Resolution

kazuma mochizuki

June 16, 2026

Abstract

The Supply Chain Management League (SCML) OneShot track requires autonomous agents to navigate complex, concurrent negotiations under strict time and computational constraints. In this report, we detail the architecture of the *Oneshot* agent, an advanced negotiation strategy specifically engineered for the ANAC 2026 competition. The proposed agent maximizes utility by combining a time-dependent Boulware pricing strategy with a highly adaptive partner selection mechanism. During early simulation stages, it adopts a broad, exploratory ordering approach to map market responsiveness, dynamically transitioning to a concentrated, exploitation-based strategy targeting trusted partners in the latter half of the simulation. Furthermore, to address the computational bottleneck inherent in combinatorial offer resolution, *Oneshot* implements a novel hybrid algorithm. By seamlessly switching between an exhaustive powerset search for small environments ($N \leq 10$) and a fast, price-sorted Greedy algorithm for larger partner sets, the agent completely eliminates timeout errors while preserving near-optimal offer matching. We demonstrate how this methodology optimizes profit margins and effectively manages quantity mismatches.

1 Introduction

The negotiation landscape in the Supply Chain Management League (SCML) OneShot track is characterized by extreme concurrency and uncertainty. Agents represent factories that must simultaneously secure raw materials from suppliers and sell finished products to consumers. Success in this environment dictates that agents not only negotiate favorable prices but also accurately balance their input and output quantities to avoid exorbitant storage costs or short-fall penalties.

Standard heuristic agents often struggle with the combinatorial explosion of concurrent offers. When an agent receives simultaneous proposals from numerous potential partners, selecting the optimal combination of offers that fulfills the exact required quantity is an NP-hard problem. Furthermore, agents who statically allocate their orders risk catastrophic failure if their chosen partners bankrupt or prove unreliable.

To overcome these challenges, we introduce the *Oneshot* agent. The core philosophy of this agent is “Scalable Cautiousness.” It employs a multi-faceted approach to negotiation, utilizing a dynamic partner selection phase, an aggressive yet adaptive pricing mechanism, and

a uniquely robust hybrid algorithm for offer resolution. This report outlines the mathematical and algorithmic foundations of the *Oneshot* agent, providing a comprehensive analysis of its internal mechanics.

2 Dynamic Partner Selection and Order Distribution

The environment within SCML is highly dynamic; factories can go bankrupt, and some agents may consistently reject offers, proving to be unreliable partners. The *Oneshot* agent addresses this by dividing the simulation into two distinct phases: Exploration and Exploitation.

2.1 Phase 1: Exploration and Market Mapping

In the first half of the simulation (defined as the current step being less than or equal to 50% of the total simulation steps), the agent aims to evaluate the responsiveness of the market. It distributes its required orders across all available, non-bankrupt partners. During this phase, a calculated *overordering fraction* is applied. By slightly over-ordering, the agent gathers crucial data on which partners are most likely to accept contracts and deliver on time.

2.2 Phase 2: Exploitation and Concentration

Once the simulation passes its midpoint, the agent's behavior shifts dramatically towards caution and reliability. The agent reviews its historical ledger ('total_agreed_quantity') and identifies partners with whom it has successfully traded.

- **Pruning:** Any partner with whom the agent has zero successful transactions after the midpoint is permanently blacklisted.
- **Concentration:** Instead of evenly distributing orders, the agent identifies its single most trusted, high-volume trading partner. The bulk of its orders is heavily concentrated onto this top-performing partner, ensuring that crucial supply chains remain intact while minimizing the risk of last-minute rejections.

3 Time-Dependent Pricing: The Boulware Strategy

The *Oneshot* agent utilizes a time-dependent concession strategy, specifically a variation of the Boulware tactic. A Boulware agent initially maintains its reservation value, refusing to concede until the negotiation deadline is imminent.

Let $t \in [0, 1]$ represent the normalized relative time within a negotiation step, where $t = 0$ is the start and $t = 1$ is the deadline. The concession rate $C(t)$ is defined exponentially:

$$C(t) = t^\epsilon \tag{1}$$

where ϵ is the concession exponent. In the *Oneshot* implementation, ϵ is set to 3.0, heavily skewing the curve so that concession remains near zero until $t > 0.7$.

Based on the allowed issue ranges for unit price, bounded by a minimum (P_{min}) and maximum (P_{max}), the price offered or countered at time t is calculated dynamically depending on the agent's role.

For the **Seller** role, the agent desires the highest possible price initially, conceding towards the minimum acceptable price:

$$P_S(t) = P_{max} - (P_{max} - P_{min}) \cdot C(t) \quad (2)$$

For the **Buyer** role, the agent desires the lowest possible price initially, conceding towards the maximum acceptable price:

$$P_B(t) = P_{min} + (P_{max} - P_{min}) \cdot C(t) \quad (3)$$

This aggressive pricing strategy ensures that the agent maximizes its profit margins against weaker negotiators early in the step, while still securing necessary contracts via rapid concession in the final milliseconds.

4 Scalable Offer Resolution: The Hybrid Algorithm

The most significant technical contribution of the *Oneshot* agent is its resolution mechanism. During the ‘counter_all’ phase, the agent must evaluate incoming offers and decide which combination of acceptances will meet its needed quantity (Q_{need}) while maximizing price utility.

4.1 The Combinatorial Explosion Challenge

If an agent has N active partners proposing offers, there are 2^N possible subsets of offers to accept. For small numbers of partners, an exhaustive search guarantees the optimal combination. However, if $N = 20$, the agent must evaluate over 1,000,000 combinations. In the time-constrained SCML environment, this $O(2^N)$ complexity inevitably leads to timeout penalties, causing the agent to forfeit the negotiation step entirely.

4.2 The Hybrid Implementation

To guarantee computational stability without sacrificing optimality, *Oneshot* implements a threshold-based hybrid algorithm:

Method A: Exhaustive Powerset Search (Active Partners ≤ 10)

When dealing with a manageable number of partners, the agent generates the full powerset of all active offers. It iterates backward through this list, tracking the combination that best minimizes the quantity difference ($Q_{offered} - Q_{need}$). If multiple combinations yield the same quantity delta, the algorithm breaks the tie by selecting the combination with the most favorable price sum (highest for sellers, lowest for buyers).

Method B: Price-Sorted Greedy Selection (Active Partners > 10)

When the number of partners exceeds the safety threshold of 10, the agent completely aborts the exhaustive search to prevent a timeout. Instead, it utilizes an $O(N \log N)$ Greedy algorithm:

1. Sort all incoming offers based on unit price. For selling, sort descending (highest price first). For buying, sort ascending (lowest price first).
2. Initialize an empty set of accepted partners and an accumulated quantity variable ($Q_{acc} = 0$).

3. Iterate through the sorted list. Add the partner’s offer to the accepted set and add their quantity to Q_{acc} .
4. Terminate the loop as soon as $Q_{acc} \geq Q_{need}$.

This hybrid approach ensures that the agent always responds within the allocated time limit, seamlessly scaling from small, intimate supply chains to massive, volatile markets.

5 Dynamic Mismatch and Over-ordering Management

Exact quantity matching is rare. Therefore, the agent employs mathematical bounds for acceptable quantity mismatches, which scale dynamically with time. The over-ordering fraction $O(t)$ dictates how much excess material the agent initially requests. It decays over time:

$$O(t) = O_{max} - (O_{max} - O_{min}) \cdot t^\gamma \quad (4)$$

Where γ (‘overordering_exp’) controls the decay curve. Similarly, the allowed mismatch tolerances ($\Delta_{min}, \Delta_{max}$) widen as the deadline approaches. If an offer combination falls outside these dynamic mismatch bounds, the agent will reject the offers and issue counter-proposals based on the Boulware pricing strategy calculated for that specific timestep.

6 Conclusion