

Strategy of SBD

SBD Team
ANAC 2026 SCML Standard Track Submission

Abstract

SBD is a rule-based agent submitted to the SCML Standard track. The agent is designed around position-dependent quantity planning, synchronized acceptance of concurrent offers, partner-aware allocation, adaptive pricing, and short-horizon future contracts. It uses a different quantity policy at the first, middle, and last production levels because these positions face different risks and opportunities. SBD does not use a large language model, deep learning, or reinforcement learning at runtime. Its behavior is produced by transparent heuristics tuned through simulation against past strong agents and standard sample agents.

1 Introduction

The SCML Standard track is a multi-agent supply-chain game in which each agent must negotiate simultaneously with several upstream and downstream partners. Each contract specifies a quantity, a delivery time, and a unit price. The agent must therefore solve two coupled problems: choosing profitable bilateral offers and coordinating all accepted contracts so that the whole set of agreements is useful under production-capacity and timing constraints.

SBD follows a compact heuristic approach. Rather than learning a policy during the tournament, it uses explicit rules for quantity targets, offer acceptance, offer generation, and concession. The central design idea is that the best quantity plan depends strongly on the agent's supply-chain position. A first level agent mainly needs to sell current supply, a last level agent mainly needs to procure enough input for exogenous sales, and a middle level agent must balance both sides while estimating whether downstream partners are actually willing to buy.

2 The Design of SBD

2.1 Negotiation Choices

At the beginning of the game, SBD identifies whether it is located at the first, middle, or last production level. This role determines the main daily quantity target.

At the first level, SBD focuses on selling available daily supply. It computes a safe target for same-day sales and avoids accepting combinations that exceed the useful current-day range. The agent evaluates candidate subsets of offers jointly, then chooses the subset that best matches the daily target while preserving utility. When the ideal quantity is not available, SBD compares the cost of underselling with the cost of taking a less balanced set of offers and selects the better alternative.

At the middle level, SBD estimates downstream demand from public exogenous summaries and from observed agreements. It avoids maintaining a fixed production-rate target independent of the market. Instead, it uses a conservative procurement cap when downstream demand is weak and increases its operating target when downstream demand appears strong. This lets the agent exploit partners that continue to buy while avoiding excessive procurement in low-demand situations.

At the last level, SBD plans purchases according to exogenous sales demand. On low-demand days it uses the realized demand plus a small safety margin. On high-demand days it switches to a capacity-based target. This hybrid policy performed better than a single fixed target because the last level faces different incentives on small-demand and large-demand days.

Prices are also role-dependent. When selling, SBD begins with favorable prices and concedes as the negotiation progresses or as quantity pressure increases. When buying, it starts with lower prices and relaxes them only when procurement is still needed. Future contracts are kept short-horizon so that they support the current plan without dominating later decisions.

2.2 Concurrent Negotiation

SBD treats same-day offers as a synchronized selection problem. In each round, the agent collects valid offers from all partners on a market side and evaluates subsets of those offers. This is important because the value of accepting one offer depends on which other offers are also accepted on the same day. For example, two individually acceptable offers may become too large when accepted together, while a lower-price offer can still be useful if it completes a target quantity.

The selected subset is scored using quantity fit, utility, and partner statistics. SBD favors offer sets that are close to the daily target and that come from partners with useful agreement histories. The same partner statistics are used when generating counter-offers: quantities are not spread uniformly across all partners, but are biased toward partners that have recently been more likely to reach agreements.

This synchronized design is especially useful in the Standard track because several negotiations are active at once. SBD therefore avoids making each acceptance decision in isolation. It instead waits for the set of available responses in a negotiation round and chooses a coordinated subset whenever

possible.

2.3 Risk Management

SBD manages risk mainly through quantity control and short planning horizons. The first-level policy prioritizes useful current-day sales and rejects excessive combinations. The middle-level policy controls procurement according to downstream demand estimates, reducing exposure when buyers appear inactive. The last-level policy uses exogenous demand as the primary signal and adds only a small margin when demand is low.

The agent also limits the influence of future contracts. It accepts and proposes future agreements only when they are close enough to current planning needs. This prevents early future commitments from forcing poor later decisions. Finally, partner reliability is tracked over time, allowing SBD to allocate more quantity to counterparts that have shown a higher probability of agreement.

3 Evaluation

SBD was evaluated in mixed Standard-track tournaments. The comparison set included strong past agents and standard sample agents, including AS0, PenguinAgent, CautiousStdAgent, XenoSotaAgent, GreedyStdAgent, RandomStdAgent, SingleAgreementAspirationAgent, and SyncRandomStdAgent. The final validation used 300 generated worlds. Scores were inspected both as an overall balanced score and separately by supply-chain position.

Agent	Balanced	First	Middle	Last
SBD final	0.969	0.414	1.154	1.338
Submitted baseline	0.954	0.380	1.135	1.347
AS0 base	0.938	0.370	1.123	1.321
PenguinAgent	0.917	0.272	1.137	1.342

Table 1: Summary of internal mixed-world evaluation. The balanced score is the average of the first-, middle-, and last-level mean scores.

The final SBD version achieved the highest balanced score in this evaluation. The largest gains appeared at the first and middle levels. At the last level, SBD remained close to the strongest baselines, while the combined result across positions was better than the reference agents used in the final validation.

4 Lessons and Suggestions

The main lesson from SBD is that a single global productivity target is not adequate for all supply-chain positions. First-level, middle-level, and last-level agents face different constraints and should therefore use different quantity policies. A second lesson is that concurrent negotiations should be handled jointly. Selecting a subset of offers is often better than accepting or rejecting each offer independently.

The design also suggests that simple partner statistics can be useful without a large predictive model. Agreement history gives enough information to bias quantity allocation toward promising partners, while the overall policy remains transparent and robust. Future work could further refine the partner model and adapt the safety margins automatically to the observed market composition.

Conclusions

SBD is a transparent rule-based agent for the SCML Standard track. It combines role-dependent quantity targets, synchronized offer selection, partner-aware allocation, adaptive prices, and short-horizon future contracts. In internal mixed-world evaluation, this design improved the balanced score against strong past agents and sample agents while keeping the strategy simple enough to inspect and maintain.