

SuperimagentZ: An agent submitted to the ANAC 2026 SCM league

Aoi Imazeki¹ Takanobu Otsuka²

^{1,2}Nagoya Institute of Technology, Aichi, Japan

¹`a.imazeki.865@stn.nitech.ac.jp`

²`otsuka.takanobu@nitech.ac.jp`

June 23, 2026

Abstract

SuperimagentZ is a heuristic agent for the SCML Standard track designed around a single objective: *keep the worst case out of ruin*. Because the score is an aggregate over many worlds, one catastrophic world—a large shortfall penalty or a pile of unsold inventory—can sink an otherwise good agent. SuperimagentZ therefore sizes every day’s buying and selling from a precise per-step account of what it can produce and is contracted to deliver, schedules future contracts with a near-future bias, prices through a time-based concession curve, and bounds every commitment with hard quantity and price caps. Evaluated in the official world distribution against both the standard built-in agents and a field of past ANAC finalists, it ranks first on mean and minimum with the lowest variance and sits within a statistical tie for the best median, never falling into a catastrophic world. In a paired, same-world comparison it is significantly ahead of every finalist except AS0, with which it ties.

1 Introduction

In the SCML Standard game an agent negotiates supply and sales contracts—possibly for future delivery—while production is capped at n_{lines} units per day. Its score in a world is the final balance normalised by the initial balance ($1.0 = \text{broke even}$), and the competition aggregates this over many randomly generated worlds. A single world in which the agent over-commits can incur a large shortfall penalty (promising more than it can deliver) or a heavy disposal/storage cost (buying more than it can sell), and such a world drags down the aggregate disproportionately. SuperimagentZ is built on the principle that follows from this: match negotiated quantities to what the agent can actually produce and is actually obliged to deliver, and never commit beyond that. Sections 2 and 3 describe the strategy and evaluate it.

2 The Design of SuperimagentZ

2.1 Need estimation

The agent extends `StdSyncAgent` and applies one uniform policy to all production levels. For each delivery step t it computes how much it still needs to buy and to sell from a

per-step inventory account rather than from aggregate needs:

$$\begin{aligned}\text{buy}(t) &= p n_{\text{lines}} - I_{\text{in}} - \text{contracted_supplies}(t), \\ \text{sell}(t) &= \min(n_{\text{lines}}, p n_{\text{lines}} + I_{\text{in}}) - \text{contracted_sales}(t),\end{aligned}$$

where $p = 0.7$ is a target productivity and I_{in} is current input inventory. Subtracting the quantities already contracted at t stops the agent from double-counting, which is the main reason it neither over-buys nor over-sells.

2.2 Exogenous-aware edge handling

At the first and last production levels one side of the chain is driven by exogenous contracts, and the agent treats those levels specially. At the first level, where raw material arrives exogenously, it does not negotiate to buy more input (buy target ≈ 0), avoiding pointless purchases that would only become disposal cost. At the last level, where the finished product is sold exogenously, it raises its buy target to the size of that exogenous sales obligation, so that it secures enough input to fulfil what it is forced to deliver and avoids shortfalls.

2.3 Negotiation, pricing and risk control

All partner negotiations run concurrently. On its own turn the agent distributes its target quantities over partners ranked by a running success rate, asking its best price. When responding, it accepts the *subset* of current offers whose total quantity is closest to the remaining need, breaking ties toward better prices and more reliable partners; offers for future delivery are spread over the next three days with a near-future bias. Prices follow a time-based concession curve: the agent opens firm and concedes toward a safe floor as the negotiation deadline approaches.

Robustness is enforced by hard caps rather than by learning. The total quantity accepted for a step is bounded by the need plus a small buffer, so a flood of offers cannot push the agent into over-commitment; acceptance thresholds tighten or loosen with inventory pressure and the time remaining; and offers whose price would be loss-making are rejected outright. These caps are what keep the worst-case world bounded.

3 Evaluation

3.1 Setup

We evaluate with the official `WorldRunner` and `ANACStdContext`, i.e. the same world distribution used by the competition. Each world draws its horizon `n_steps` uniformly from $[20, 100]$; the reported figures therefore aggregate a *mixture* of horizons, but every world has a single fixed `n_steps`. All comparisons are *paired*: every agent under test is run in the *same* generated worlds—hence at the same `n_steps`—so both between-world difficulty and horizon-length cancel out and the differences reflect the agents themselves. We report the mean and standard deviation together with the minimum, median and interquartile range (IQR), because in this game the lower tail matters as much as the centre.

We run two studies. The first (Table 1) compares SuperimagentZ against the two standard built-in agents. The second (Table 2) places it in a field of *past ANAC finalists* loaded from `scml_agents`—the strongest finishers of the 2025 edition, AS0, XenoSota and UltraSuperMiracle, together with PenguinAgent (2024)—with `SyncRandomStdAgent` as a reference, all evaluated over the same 70 worlds. To make the environment realistic rather

than a self-built straw opponent, the background population in both studies is a mix of implementation-diverse built-in agents and the 2025 ANAC finalists (one agent per team). For the finalist study we also report the paired mean difference of SuperimagentZ minus each opponent with a 2 standard-error band.

3.2 Results

Table 1 shows that SuperimagentZ leads the built-in baselines on every statistic: a markedly higher mean and median, the smallest standard deviation, and the highest minimum. The decisive point is the minimum: SuperimagentZ’s worst world still ends in profit-neutral territory (0.258), whereas **SyncRandomStdAgent** suffers genuinely catastrophic worlds (minimum -0.742 , i.e. losing most of its starting capital). In other words, the per-step accounting and hard caps do what they were designed to do—they remove the catastrophic tail rather than merely raising the average.

Agent	Mean	Std	Min	Median	25%–75%
SuperimagentZ	1.128	0.248	0.258	1.162	1.056–1.261
SyncRandomStdAgent	0.813	0.437	-0.742	0.842	0.669–1.001
GreedyStdAgent	0.844	0.324	0.162	0.919	0.632–1.054

Table 1: Paired normalised score over the same worlds (variable horizon). SuperimagentZ keeps its minimum well clear of ruin; the random baseline does not.

The harder test is against real competitors. Table 2 places SuperimagentZ in a field of strong ANAC finalists over the same worlds. On the marginal distribution the top four agents are tied on the median within a span of 0.006, so the median alone does not separate them. What separates them is everything else: SuperimagentZ has the highest *mean* (0.969), the highest *minimum* (0.244), and the lowest dispersion on both measures (standard deviation 0.226, IQR 0.230). It is the most stable agent and the least exposed to disastrous worlds—for comparison, XenoSota reaches a slightly higher median but does so with a far worse minimum (0.015) and larger variance.

The paired comparison resolves the tie. On identical worlds, SuperimagentZ is significantly ahead of XenoSota ($+0.028 \pm 0.027$), UltraSuperMiracle ($+0.066 \pm 0.047$) and PenguinAgent ($+0.046 \pm 0.031$); against AS0 the difference is $+0.014 \pm 0.034$, a statistical tie. We state this honestly: SuperimagentZ ties AS0 and edges ahead of the rest, chiefly by holding the better worst case and the lower variance. This is also a reminder that the marginal ranking can mislead—XenoSota’s higher median comes from higher-variance worlds, and on the *same* worlds SuperimagentZ is the one ahead.

Figure 1 shows the full score distributions and makes the point visually: SuperimagentZ’s box is the most compact and its lower whisker the highest among the strong agents, while the higher-median finalists trail long lower tails into loss territory.

Conclusions

SuperimagentZ shows that, in a game where one disaster can dominate the aggregate score, sizing commitments to true production capacity and contractual obligations—and bounding them with hard caps—reliably removes the catastrophic worlds while keeping a competitive average. Its strength is this worst-case discipline as a whole rather than any single trick: in controlled ablation the exogenous-aware edge handling we expected




Figure 1: Normalised-score distribution over the same 70 worlds (boxes ordered by median; SuperimagentZ highlighted; dashed line = break-even). SuperimagentZ has the tightest interquartile box and the highest lower tail among the strong agents—the worst-case discipline shows up as a short lower whisker, not merely a high centre.

Agent	Median	Mean	Std	Min	5%	IQR
XenoSota	0.996	0.941	0.294	0.015	0.356	0.339
SuperimagentZ	0.993	0.969	0.226	0.244	0.536	0.230
UltraSuperMiracle	0.992	0.903	0.309	-0.039	0.323	0.400
AS0	0.990	0.955	0.281	0.239	0.384	0.343
PenguinAgent	0.980	0.924	0.296	0.124	0.334	0.394
SyncRandomStdAgent	0.834	0.745	0.341	-0.503	0.224	0.422

Table 2: Normalised score against strong ANAC finalists over the same 70 worlds, sorted by median. The top four tie on the median, but SuperimagentZ leads on mean, minimum and both dispersion measures (bold). In the paired same-world comparison SuperimagentZ is significantly ahead of every finalist except AS0 ($+0.014 \pm 0.034$, a tie).

to be the decisive improvement turned out to be roughly neutral on average (helping the edge-level minimum only slightly), and a number of further extensions, including a rolling-horizon planning variant, did not improve on the design either. Pushing meaningfully beyond this robustness ceiling likely calls for a different architecture—uncertainty-aware multi-day planning—which we leave to future work.