Design of an Agent with Partner Scoring and Market Trend Analysis for the SCML Standard Track

[1]Yuta Kawasaki, [2]Ryuta Shiraishi
Nagoya Institute of Technology, Aichi, Japan
[1]y.kawasaki.083@stn.nitech.ac.jp, [2]r.shiraishi.749@stn.nitech.ac.jp

June 14, 2025

Abstract

This report explains the design of PriceTrendStdAgent, which was submitted to the ANAC2025 SCML Standard Track. The agent is based on CautiousStdAgent from 2024, and adds two main ideas: scoring negotiation partners using past trades, and using price trends to decide when and how much to buy or sell. The agent calculates scores from trade prices and quantities, checks market trends with moving averages, and changes its offers depending on the situation. These improvements help the agent make better decisions and earn more profit.

1 Introduction

PriceTrendStdAgent is an enhanced version of CautiousStdAgent, which was submitted to the ANAC 2025 Standard. In addition to the cautious decision-making strategy based on inventory status used by CautiousStdAgent, this agent introduces the following three major improvements:

- Scoring partners based on the price and quantity of successful trades
- Market trend prediction using moving averages
- Dynamic adjustment of proposal timing and quantity based on market trend prediction

These enhancements aim to improve the agent' s long-term profitability and adaptability to market changes.

2 The Design of PriceTrendStdAgent

2.1 Partner Scoring

PriceTrendStdAgent records each successful transaction and updates a cumulative score S_p for each negotiation partner p. The cumulative score is calculated based on the following two components:

- Price Score : The normalized position of the negotiated price within the range between the minimum and maximum prices.
- Quantity Score : The normalized traded quantity based on the maximum daily processing capacity.

The score S_p for each negotiation partner is updated according to the following formula:

$$S_p \leftarrow \alpha(\text{price\_score} + \beta\text{quantity\_score}) + (1-\alpha)S_p$$

The parameters are set to $\alpha = 0.4$ and $\beta = 1.5$. These scores are updated using exponential smoothing, with weights of 0.6 for the past score and 0.4 for the most recent trade. In the latter half of the negotiation process, the agent switches from a uniform allocation strategy to a prioritized distribution strategy that allocates larger quantities to negotiation partners with higher overall scores S_p.

2.2 Market Trend Prediction

PriceTrendStdAgent monitors the changes in market prices and predicts price trends using exponential moving averages (EMA). First, for the time series of input prices Price_t, which represent the average market prices (i.e., the mean of executed trade prices) of input products at each step as provided by the SCML environment, the exponential moving average is updated using the following equation:

$$\text{EMA}_t = \alpha \text{Price}_t + (1 - \alpha)\text{EMA}_{t-1}$$

Here, $\alpha = 0.3$ is used. Based on the difference $\Delta = \text{Price}_t - \text{EMA}_t$ between the current price and the EMA, the trend is classified as follows:

- $\Delta > 0.5$: Upward trend
- $\Delta < -0.5$: Downward trend
- $|\Delta| \leq 0.5$: Stable

In this approach, a "trend" is identified based on whether the difference Δ exceeds a threshold of 0.5. This difference acts as a simple indicator of how significantly the current price deviates from past averages, capturing both the direction and magnitude of the trend. If $|\Delta| \leq 0.5$, the price is considered stable. Otherwise, it is regarded as trending and is used in offer decision-making. This threshold was empirically chosen to avoid overreacting to normal fluctuations in price.

2.3 Trend-aware Timing and Quantity Adjustment

When evaluating offers from negotiation partners, the agent considers the Unit price, Delivery Day, and quantity in that order. For trade price and quantity, it uses the same criteria as CautiousStdAgent: the price must be within the acceptable range, and the quantity must match current inventory levels or demand.

However, for trade timing, PriceTrendStdAgent incorporates market price trend detection to make the following strategic decisions:

- Upward trend: The agent tends to accelerate the purchase, expecting prices to rise further.
- Downward trend: The agent tends to delay the purchase by one step, expecting prices to fall further.

If an offer is scheduled for a future step but an upward trend is detected, the agent modifies the proposed time to the current step and sends a counteroffer to avoid future price increases. Conversely, under a downward trend, the agent shifts the proposed time one step later to benefit from the expected price drop.

In addition, the proposed quantity is adjusted based on the identified price trend using the following multipliers:

- Upward trend: 0.9 (decrease)
- Downward trend: 1.1 (increase)
- Stable: 1.0 (no change)

These adjustments allow the agent to avoid excessive buying or selling during price surges and instead increase volume during price drops to maximize profit. By doing so, PriceTrendStdAgent aims to both mitigate risk and exploit opportunities for profit during downward price phases.

3 Evaluation

To evaluate the effectiveness of the proposed PriceTrendStdAgent, we conducted a series of benchmark experiments against the following three agents, all of which achieved high rankings in 2024 competitions. CautiousStdAgent is the second-place agent from ANAC2024 and serves as the foundation for our proposed method. PenguinAgent and DogAgent were the first- and third-place agents, respectively, in the same competition.

In this experiment, The configurations are n_configs = 30 and n_competitors_per_world = 4. The results are shown in Table 1.

Table 1: Mean Scores of Agents in the Tournament (bold = 1st, underline = 2nd)

Agent Name	**Step 50**	**Step 100**	**Step 200**
PriceTrendStdAgent (our)	<u>1.0109</u>	<u>0.9914</u>	<u>0.8103</u>
CautiousStdAgent	1.0020	0.9803	0.8102
PenguinAgent	**1.0580**	**1.0558**	**0.8414**
DogAgent	1.0041	0.9802	0.7735

Table 2 provides a more detailed statistical results in the Step 100 setting. It includes the mean, standard deviation, minimum, 25th percentile, median, 75th percentile, and maximum score for each agent.

Table 2: Detailed Results of the Tournament (Step 100) (bold = 1st, underline = 2nd in mean/std)

Agent Name	**Mean**	**Std**	**Min**	**25%**	**Median**	**75%**	**Max**
PriceTrendStdAgent (our)	<u>0.9914</u>	<u>0.3095</u>	-0.1100	0.8602	1.0000	1.1921	1.5256
CautiousStdAgent	0.9803	**0.3018**	-0.0817	0.8537	1.0000	1.1920	1.5098
PenguinAgent	**1.0558**	0.3786	-0.2750	0.8546	1.1528	1.3385	1.6693
DogAgent	0.9802	0.3401	-0.3202	0.8005	1.0000	1.2133	1.6539

In experiments with 50- and 100-step runs, PriceTrendStdAgent achieved higher mean scores than CautiousStdAgent However, when the step count increases to 200, the performance gap between the two agents narrows and their scores tend to align more closely. These findings suggest that the proposed trend-based adjustment strategy is effective in short-term negotiation scenarios. It is possible that, as negotiations progress, market prices stabilize, thereby reducing the effectiveness of trend-based adjustments. Nevertheless, the specific reasons why trend-based strategies become less effective in longer-term scenarios cannot be fully understood from these results alone. One possible factor is that market price fluctuations may become less pronounced over time. Further observation and analysis are required to verify this hypothesis.

In all step settings, PriceTrendStdAgent was not able to outperform PenguinAgent. While both CautiousStdAgent and PriceTrendStdAgent rely on cautious, short-term decision-making, PenguinAgent uses an aggressive strategy that anticipates future steps from the proposal stage. This difference in negotiation style may explain PenguinAgent' s higher profitability. As shown in Table 2, PenguinAgent not only achieves the highest mean score but also attains the highest maximum score, confirming the effectiveness of its forward-looking strategy.

These results suggest that the proposed improvements are effective in short-term negotiations. However, additional improvements will be necessary to remain competitive in longer-term scenarios or against agents that incorporate future-oriented strategies.

4 Conclusion and Suggestions

We proposed PriceTrendStdAgent, an extension of the CautiousStdAgent, by adding two key improvements: partner scoring based on trade quantities and prices, and dynamic adjustments using market price trends.

Experimental results demonstrated that these improvements are effective in short-term negotiations, showing higher mean scores compared to the base agent.

Based on these findings, we propose the following future directions to further enhance the agent' s performance: First, implementing forward-looking negotiation strategies could enable the agent to make proposals and responses with a longer-term perspective. Second, investigating the effectiveness of the scoring mechanism and optimizing its update parameters may help better reflect partner reliability and negotiation outcomes.