

SCML Agent MMM Team

Submitted by:

Ohad Amosy
Lior Strichash
Yosef Noorani

Background

Introduction

The OneShot game is a simplification of the SCML standard game, designed to de-emphasize long-term production and planning, and instead, emphasize negotiations.

We chose to improve existing strategies and solve some of their drawbacks.

Strategy

Our approach is a semi-greedy agent that attempts to maximize profit.

The agent is an enhancement of several agents from the documentation.

One of the existing suggested OneShot agents is "BetterAgent" which is a greedy agent that uses concessions exponent.

The agent does not take into account the amounts from previous negotiations.

The agent "AdaptiveAgent" considers previous negotiations and limits the concessions.

This logic aims to improve the agent's performance over time.

However, this approach makes the agent hard-headed by not compromising the price it demands, once it gets a better price.

What did we improve in the AdaptiveAgent?

What is a good price?

We used the estimation of the AdaptiveAgent for a good price and added a profit margin to the offer we send. Another strategy we tried is to save the best prices over time (like agent 124 from last year). Instead of using the average, we used the 80th percentile. We tested for the 20, 50, 80, 90 percentiles and 80 is the best value.

Should we buy more than we need?

We assume that some agents will not be able to execute the agreement, so we tried to elevate this by buying (and agreeing to sell) more than our actual need. We tried that for different margin values [5%,10%,20%].

Can we assume all agents will meet their obligations?

All the agents from the documentation offer to buy from a single agent all the goods. They calculate “needs” and offer it to all the agents. None of them try to negotiate different amounts. We tried to diversify our partners by offering different quantities to different agents. We based our quantity logic on the best price we got from each agent. The higher the price a partner sells to us, the more quantity we are willing to sell to it. We tested two distribution functions: softmax and linear.

Personalization

We tried several ways to take into account the fact that different agents may have different strategies. For each day we saved for each agent all the offers it sends and all the signed contracts. Then we use this data to fine-tune the global price for each agent using the following:

1. Average the global price and the best contract it signed with us in the past.
2. Average the global price with the current offers other agents send us.
3. Average best sign offer we have from this agent with the current offers other agents send us.

We also tried to give more “attractive” offers (5%-10% discount) to agents that do not sign with us in the past. The main motivation was to

learn what was their threshold and use that to update our offers to other agents.

In the opposite direction, we tried to reject offers from agents in the first steps (tried 1-5 steps), to learn for each agent how it will improve an offer after rejection.