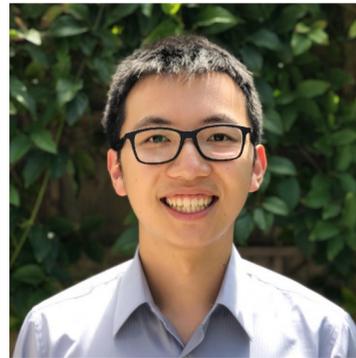


# Diverse Conventions for Human-AI Collaboration



Bidipta Sarkar



Andy Shih



Dorsa Sadigh

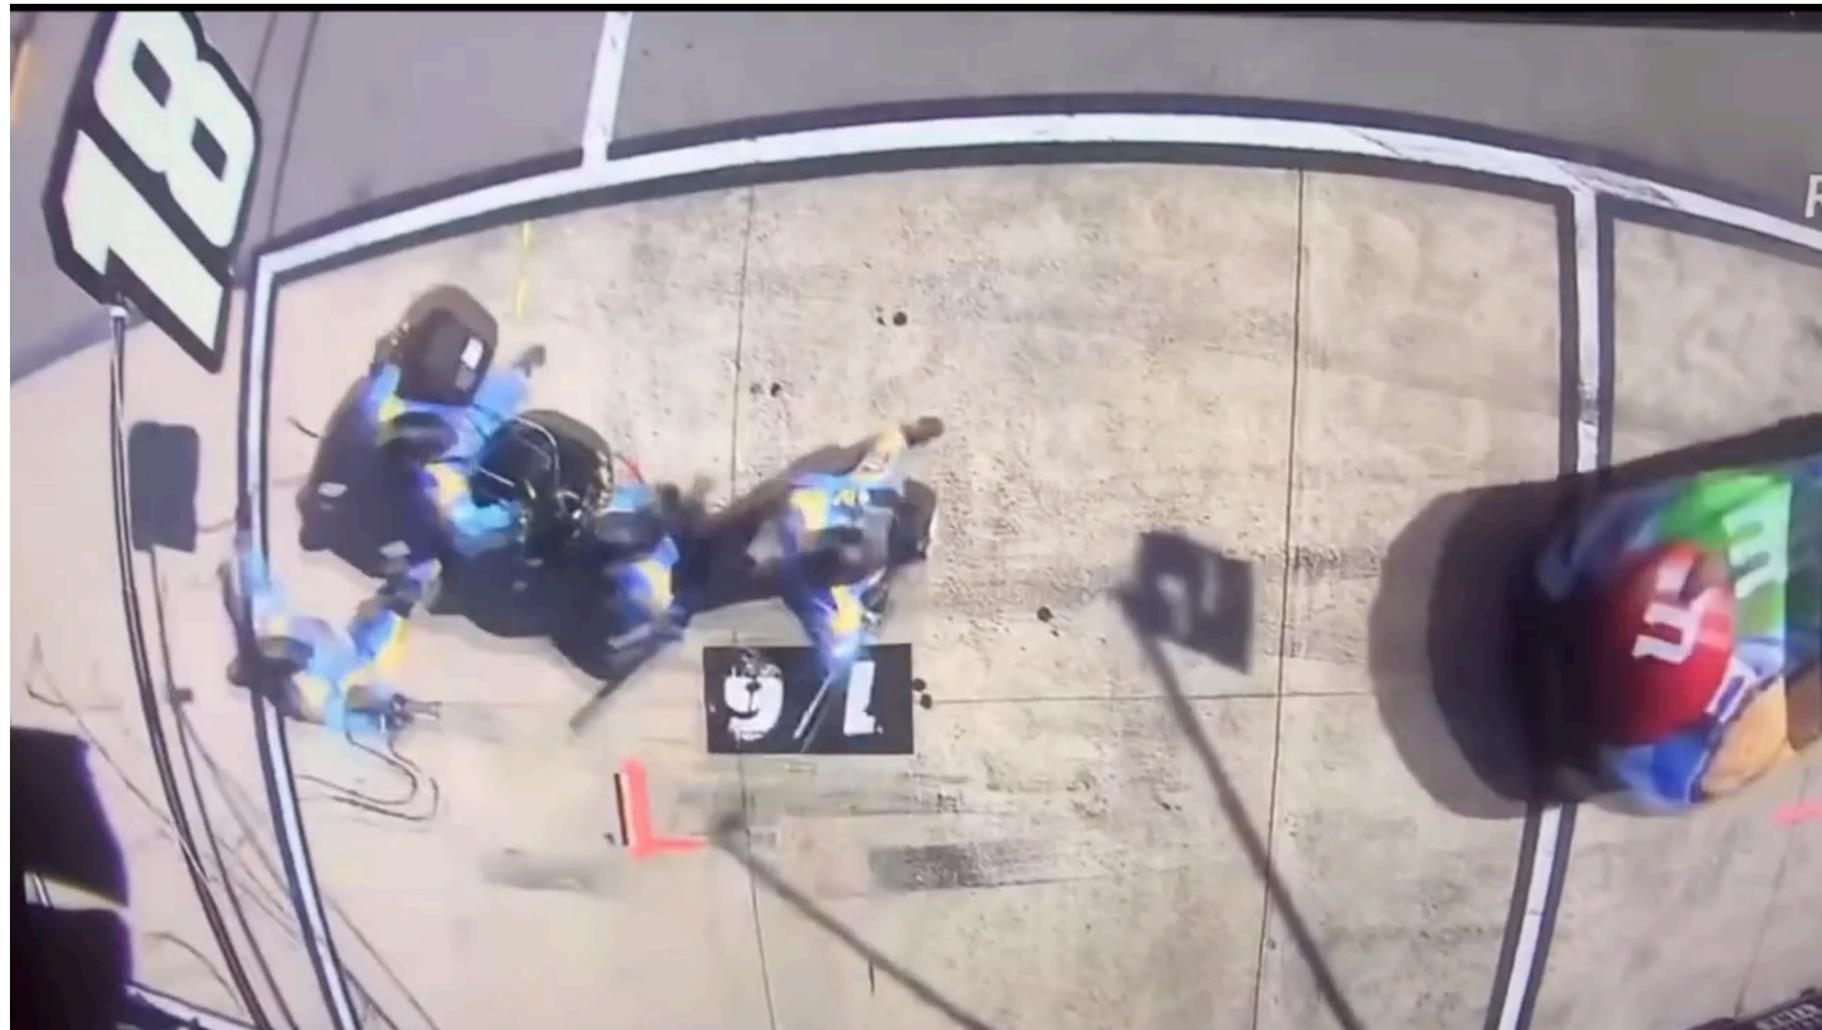


# Conventions

Arbitrary solutions to recurring coordination problems

# Conventions

Arbitrary solutions to recurring coordination problems



# AI Conventions



# AI Conventions



*Two AI Agents achieve expert human performance*

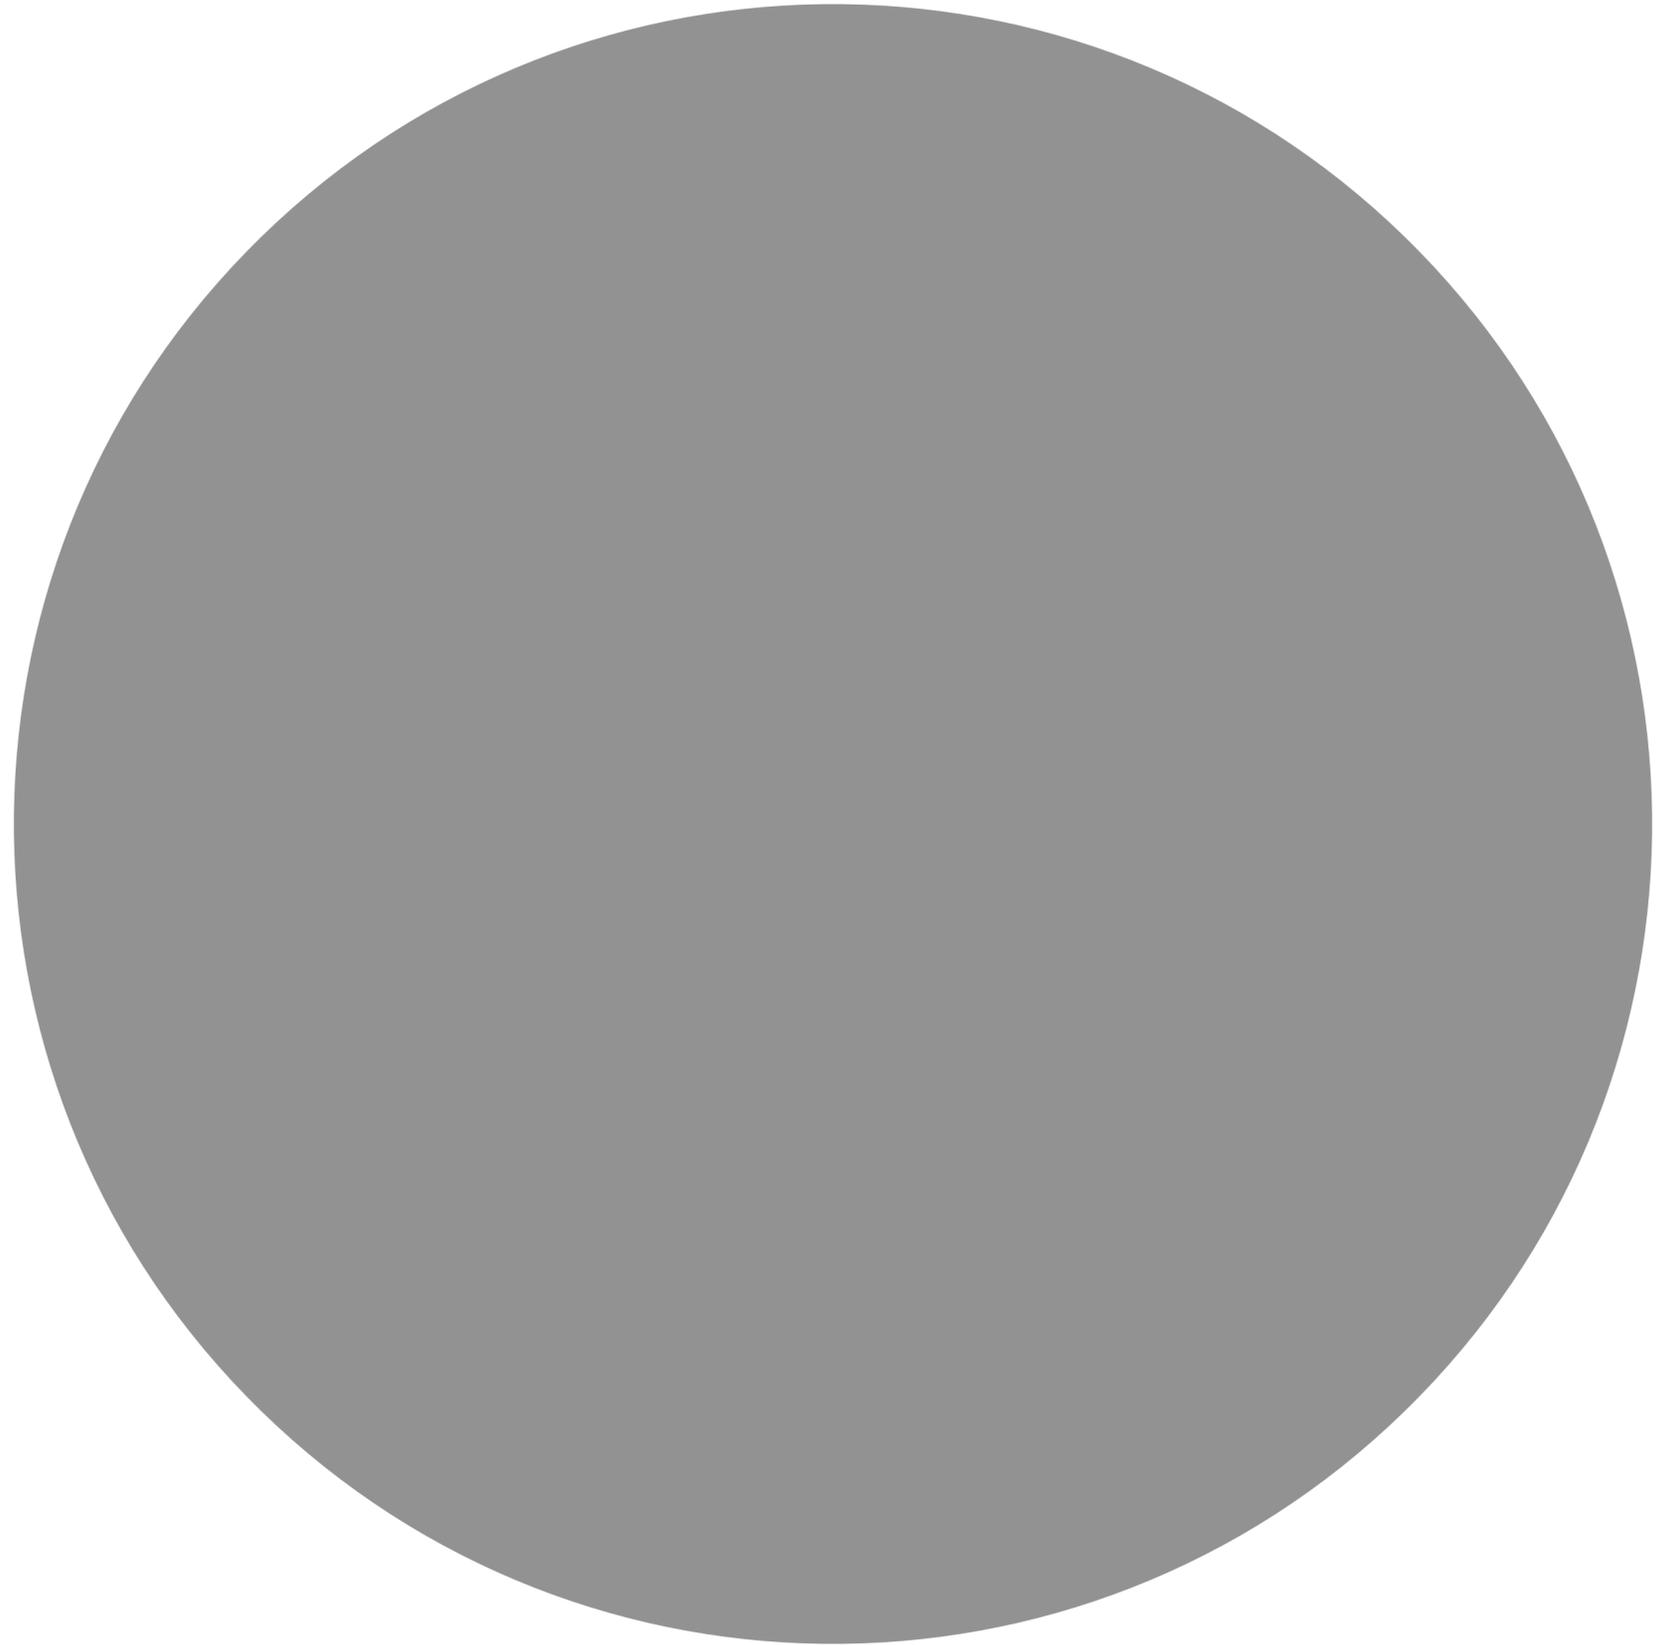
# AI Conventions

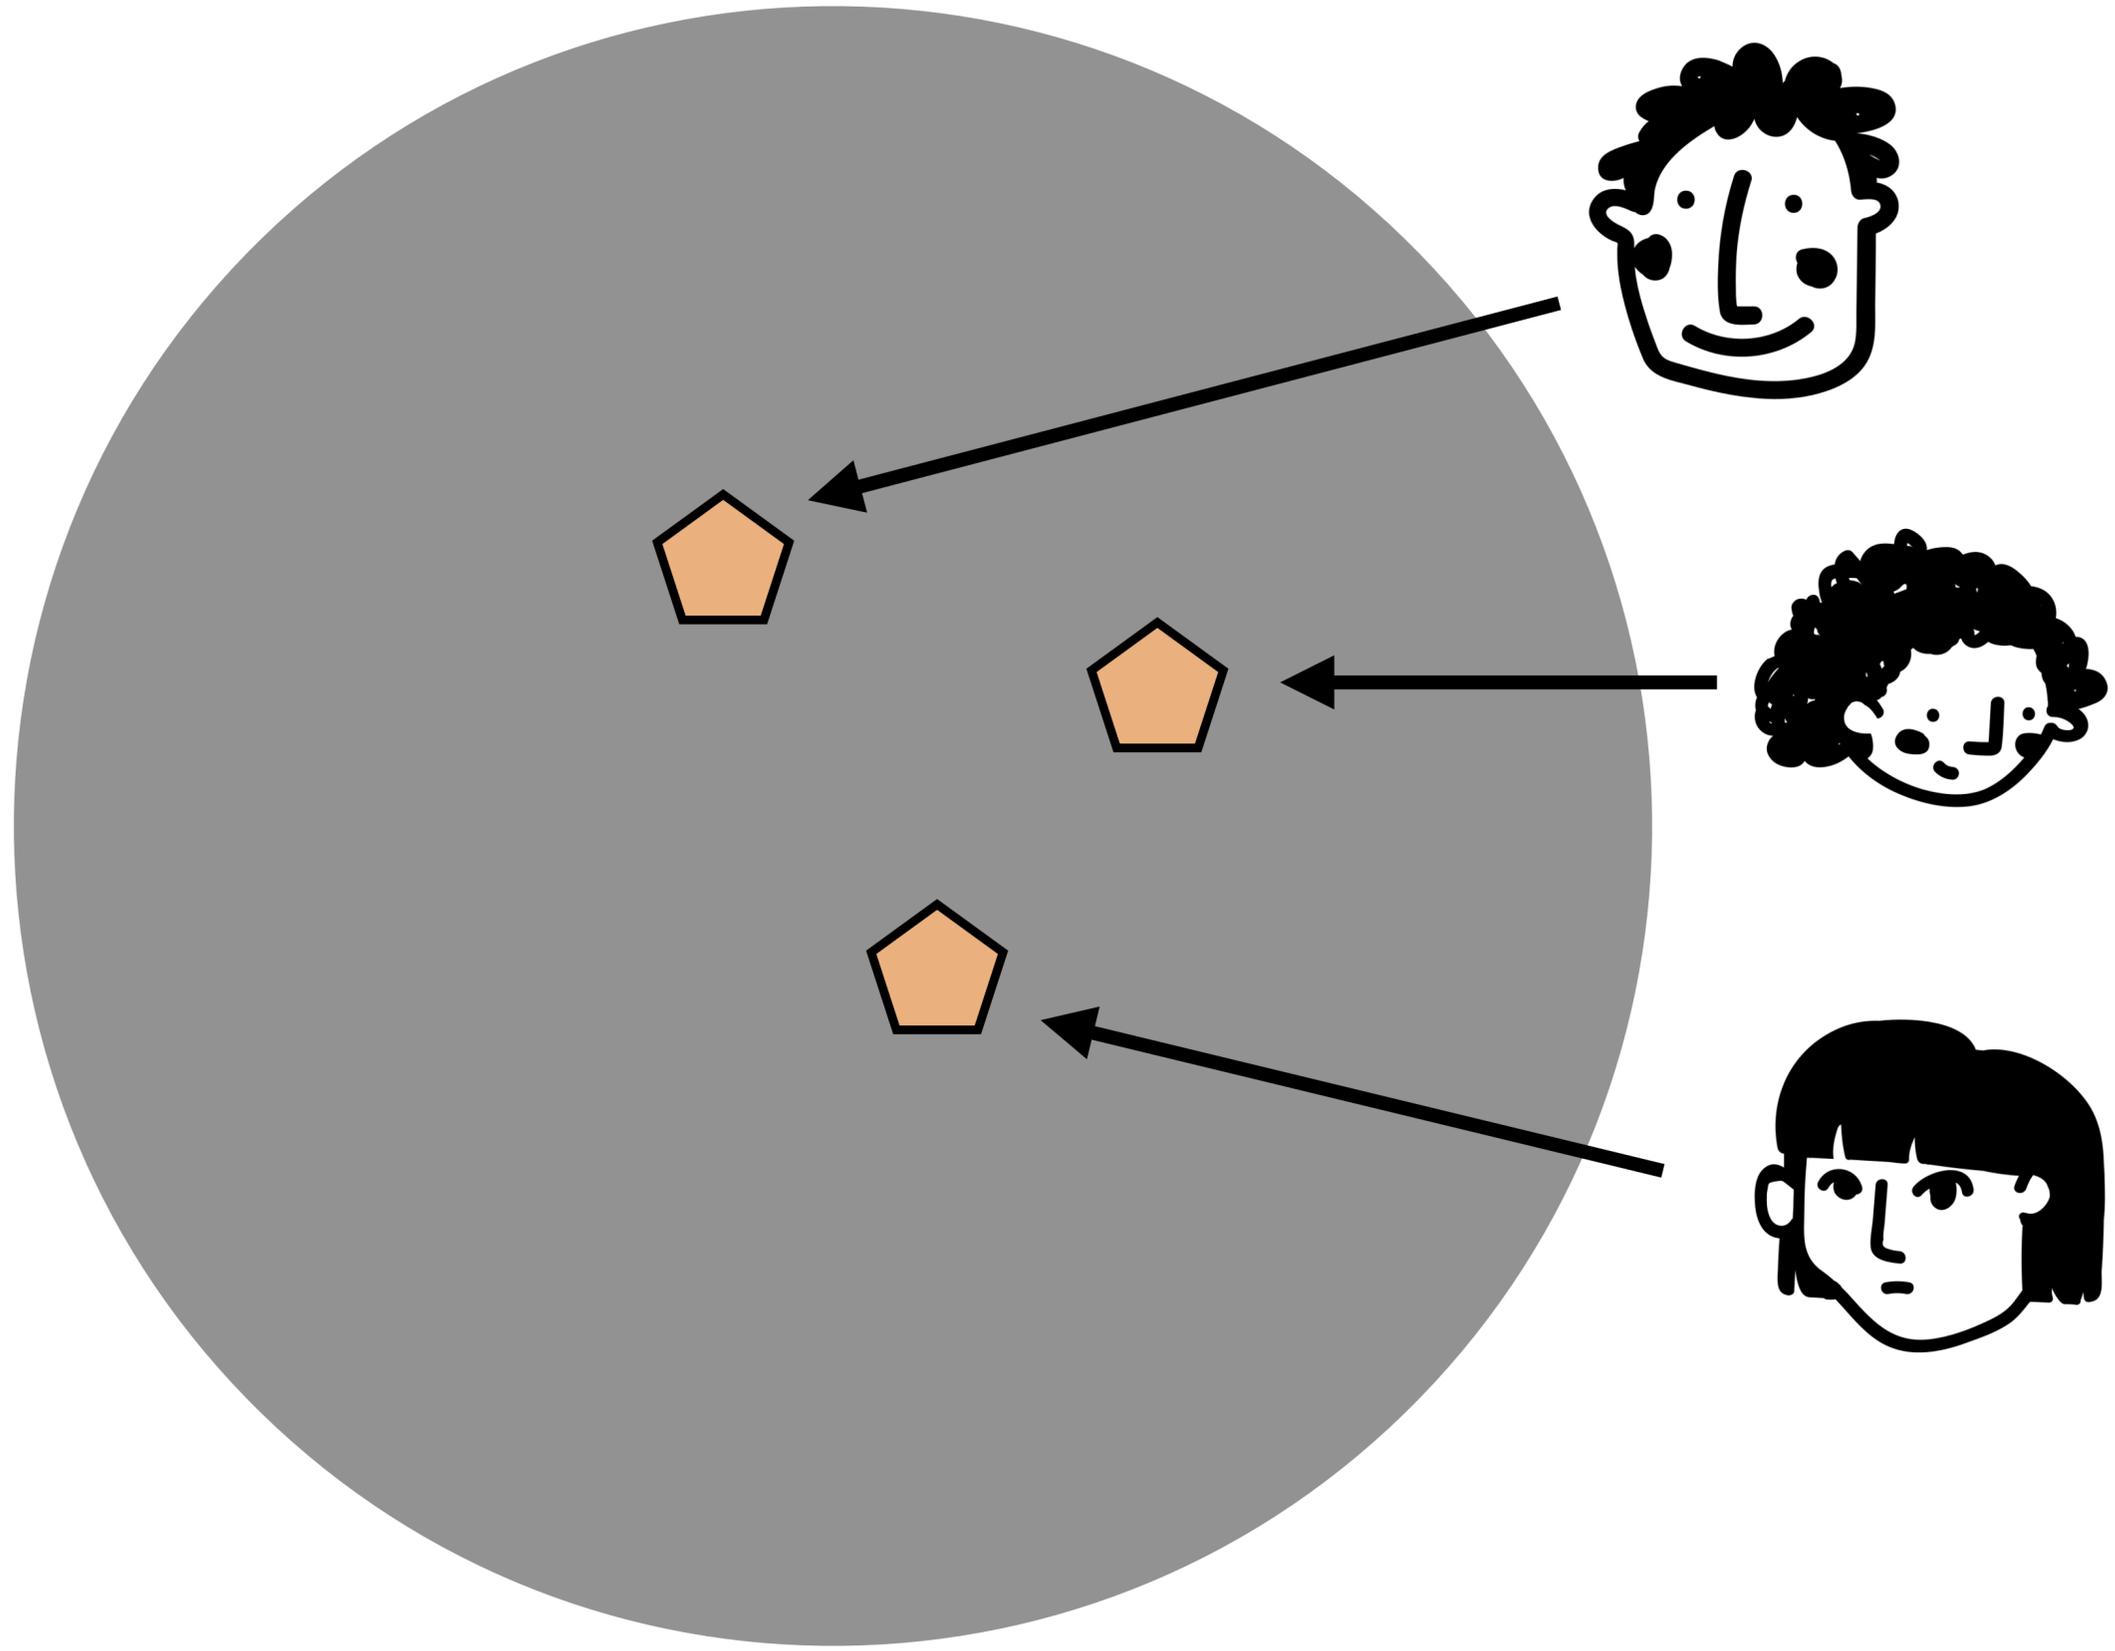


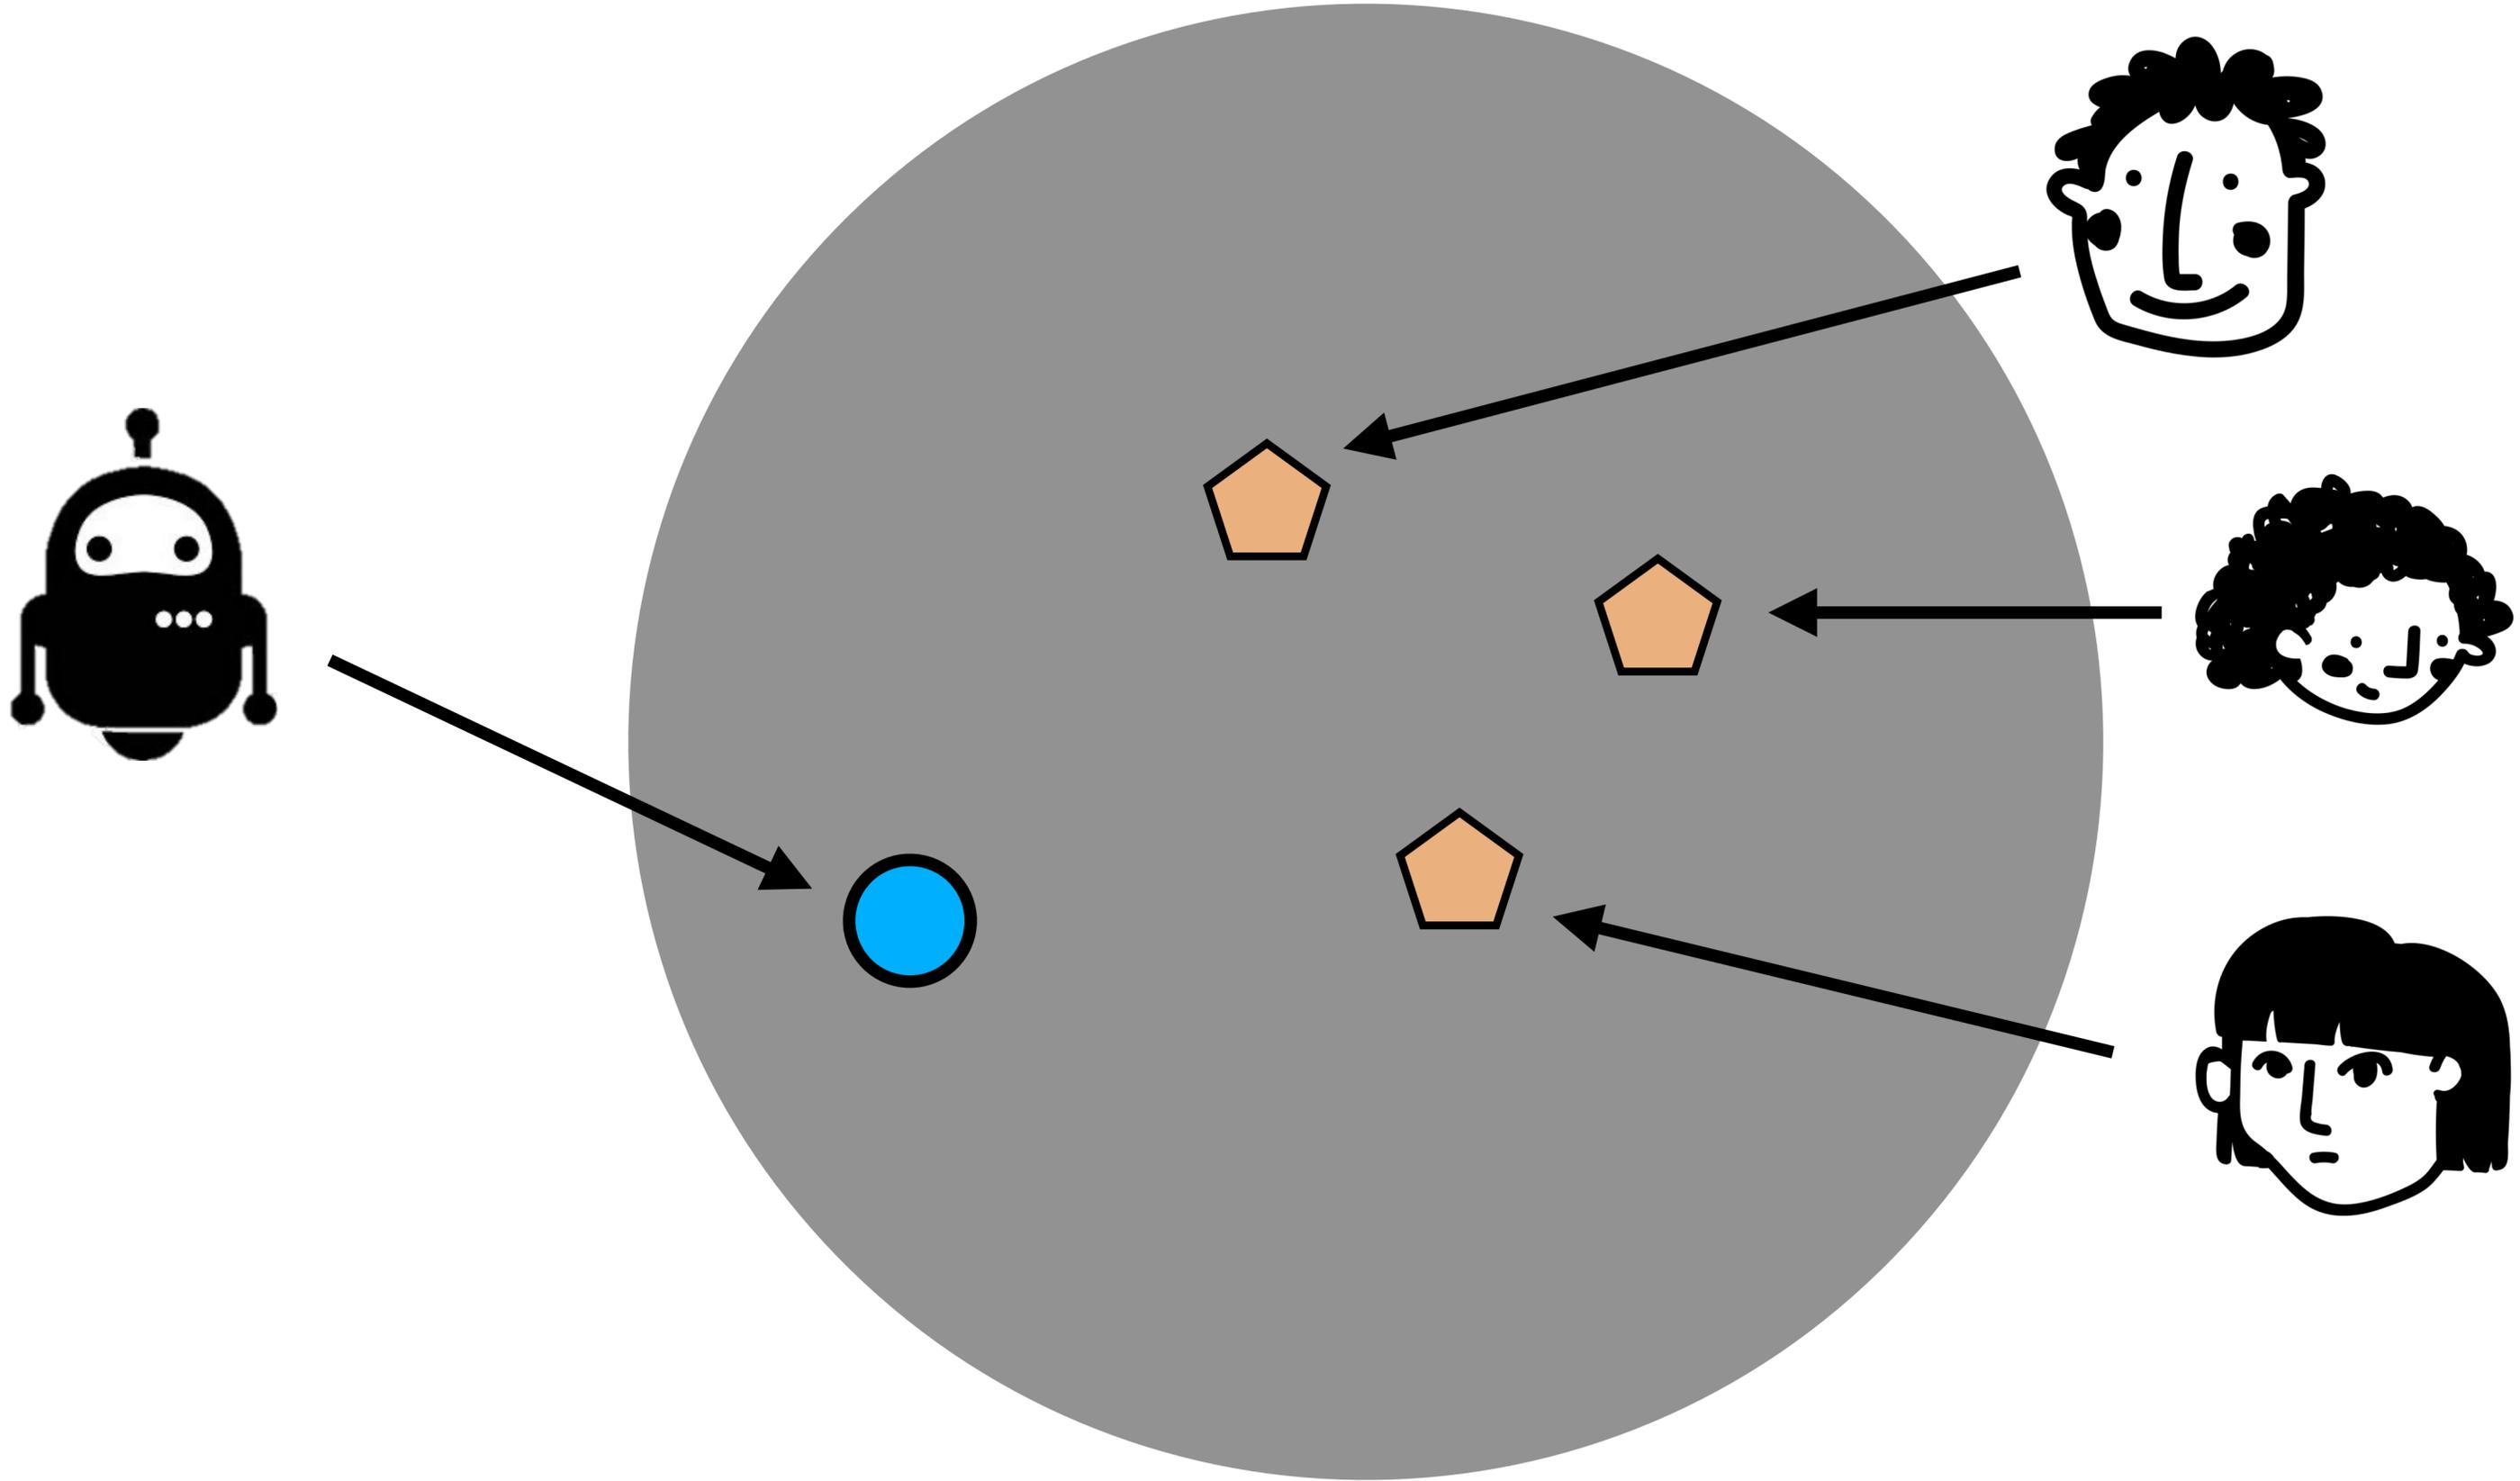
*Two AI Agents achieve expert human performance*

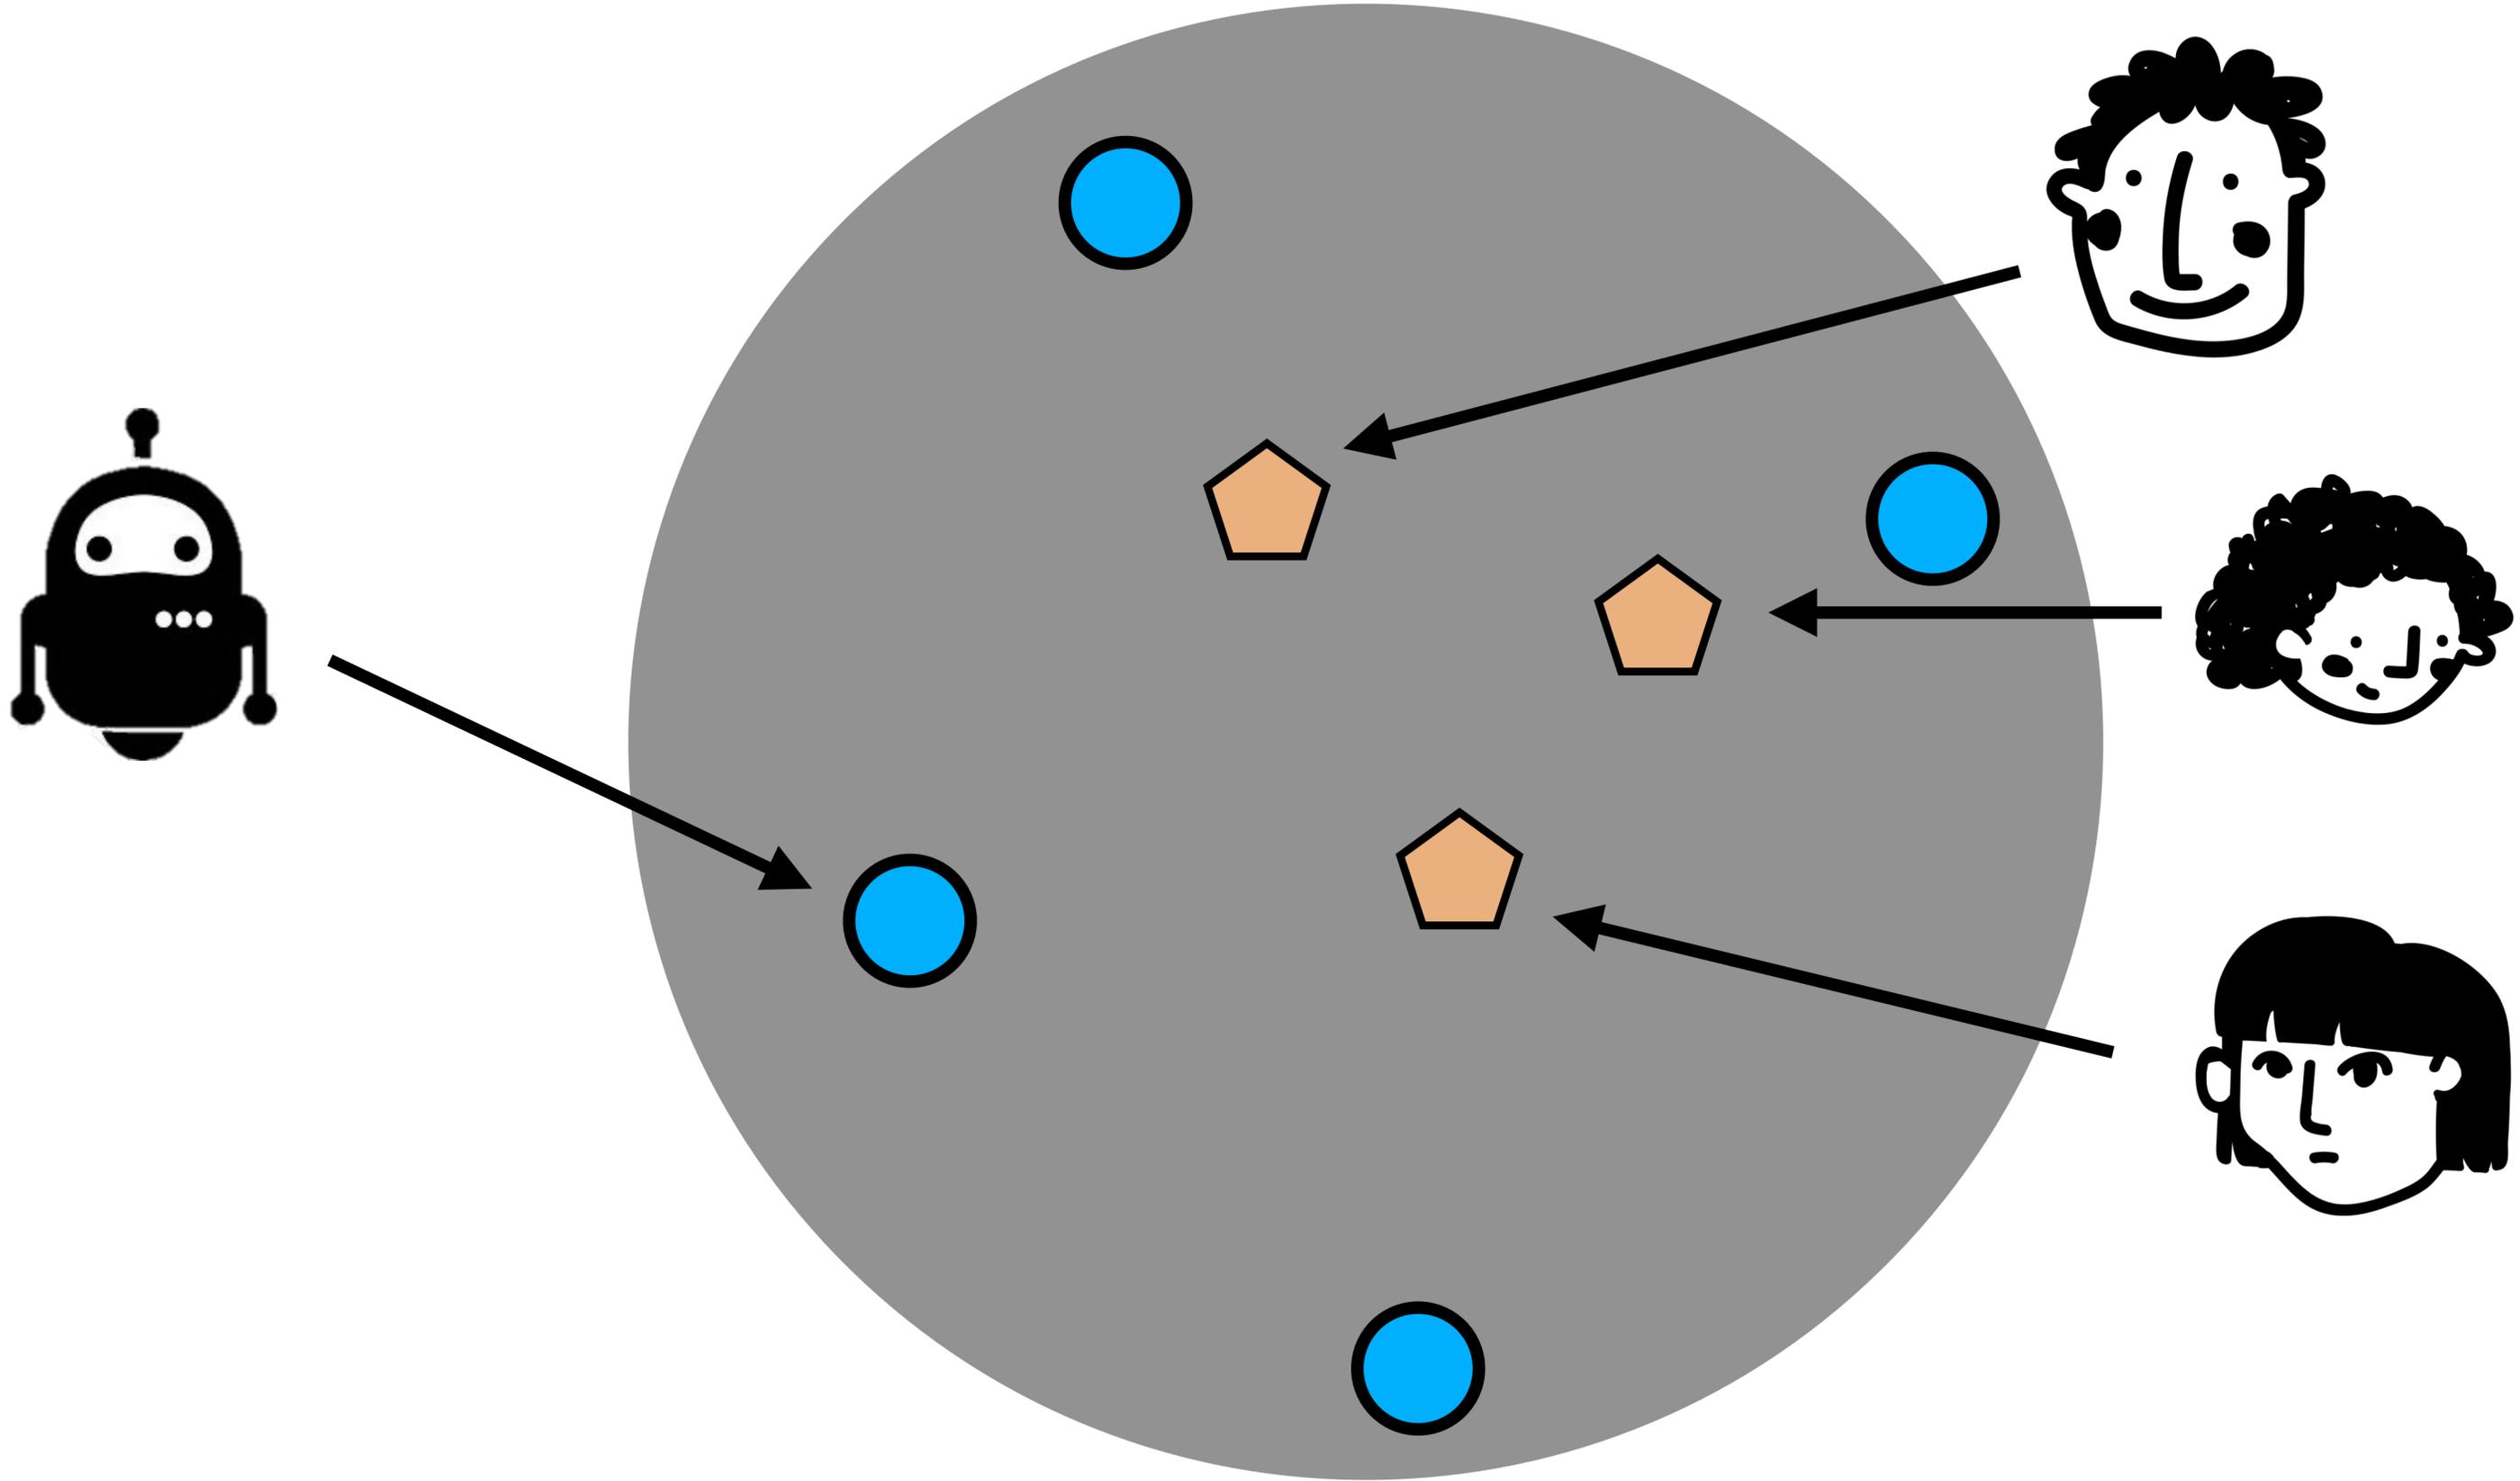


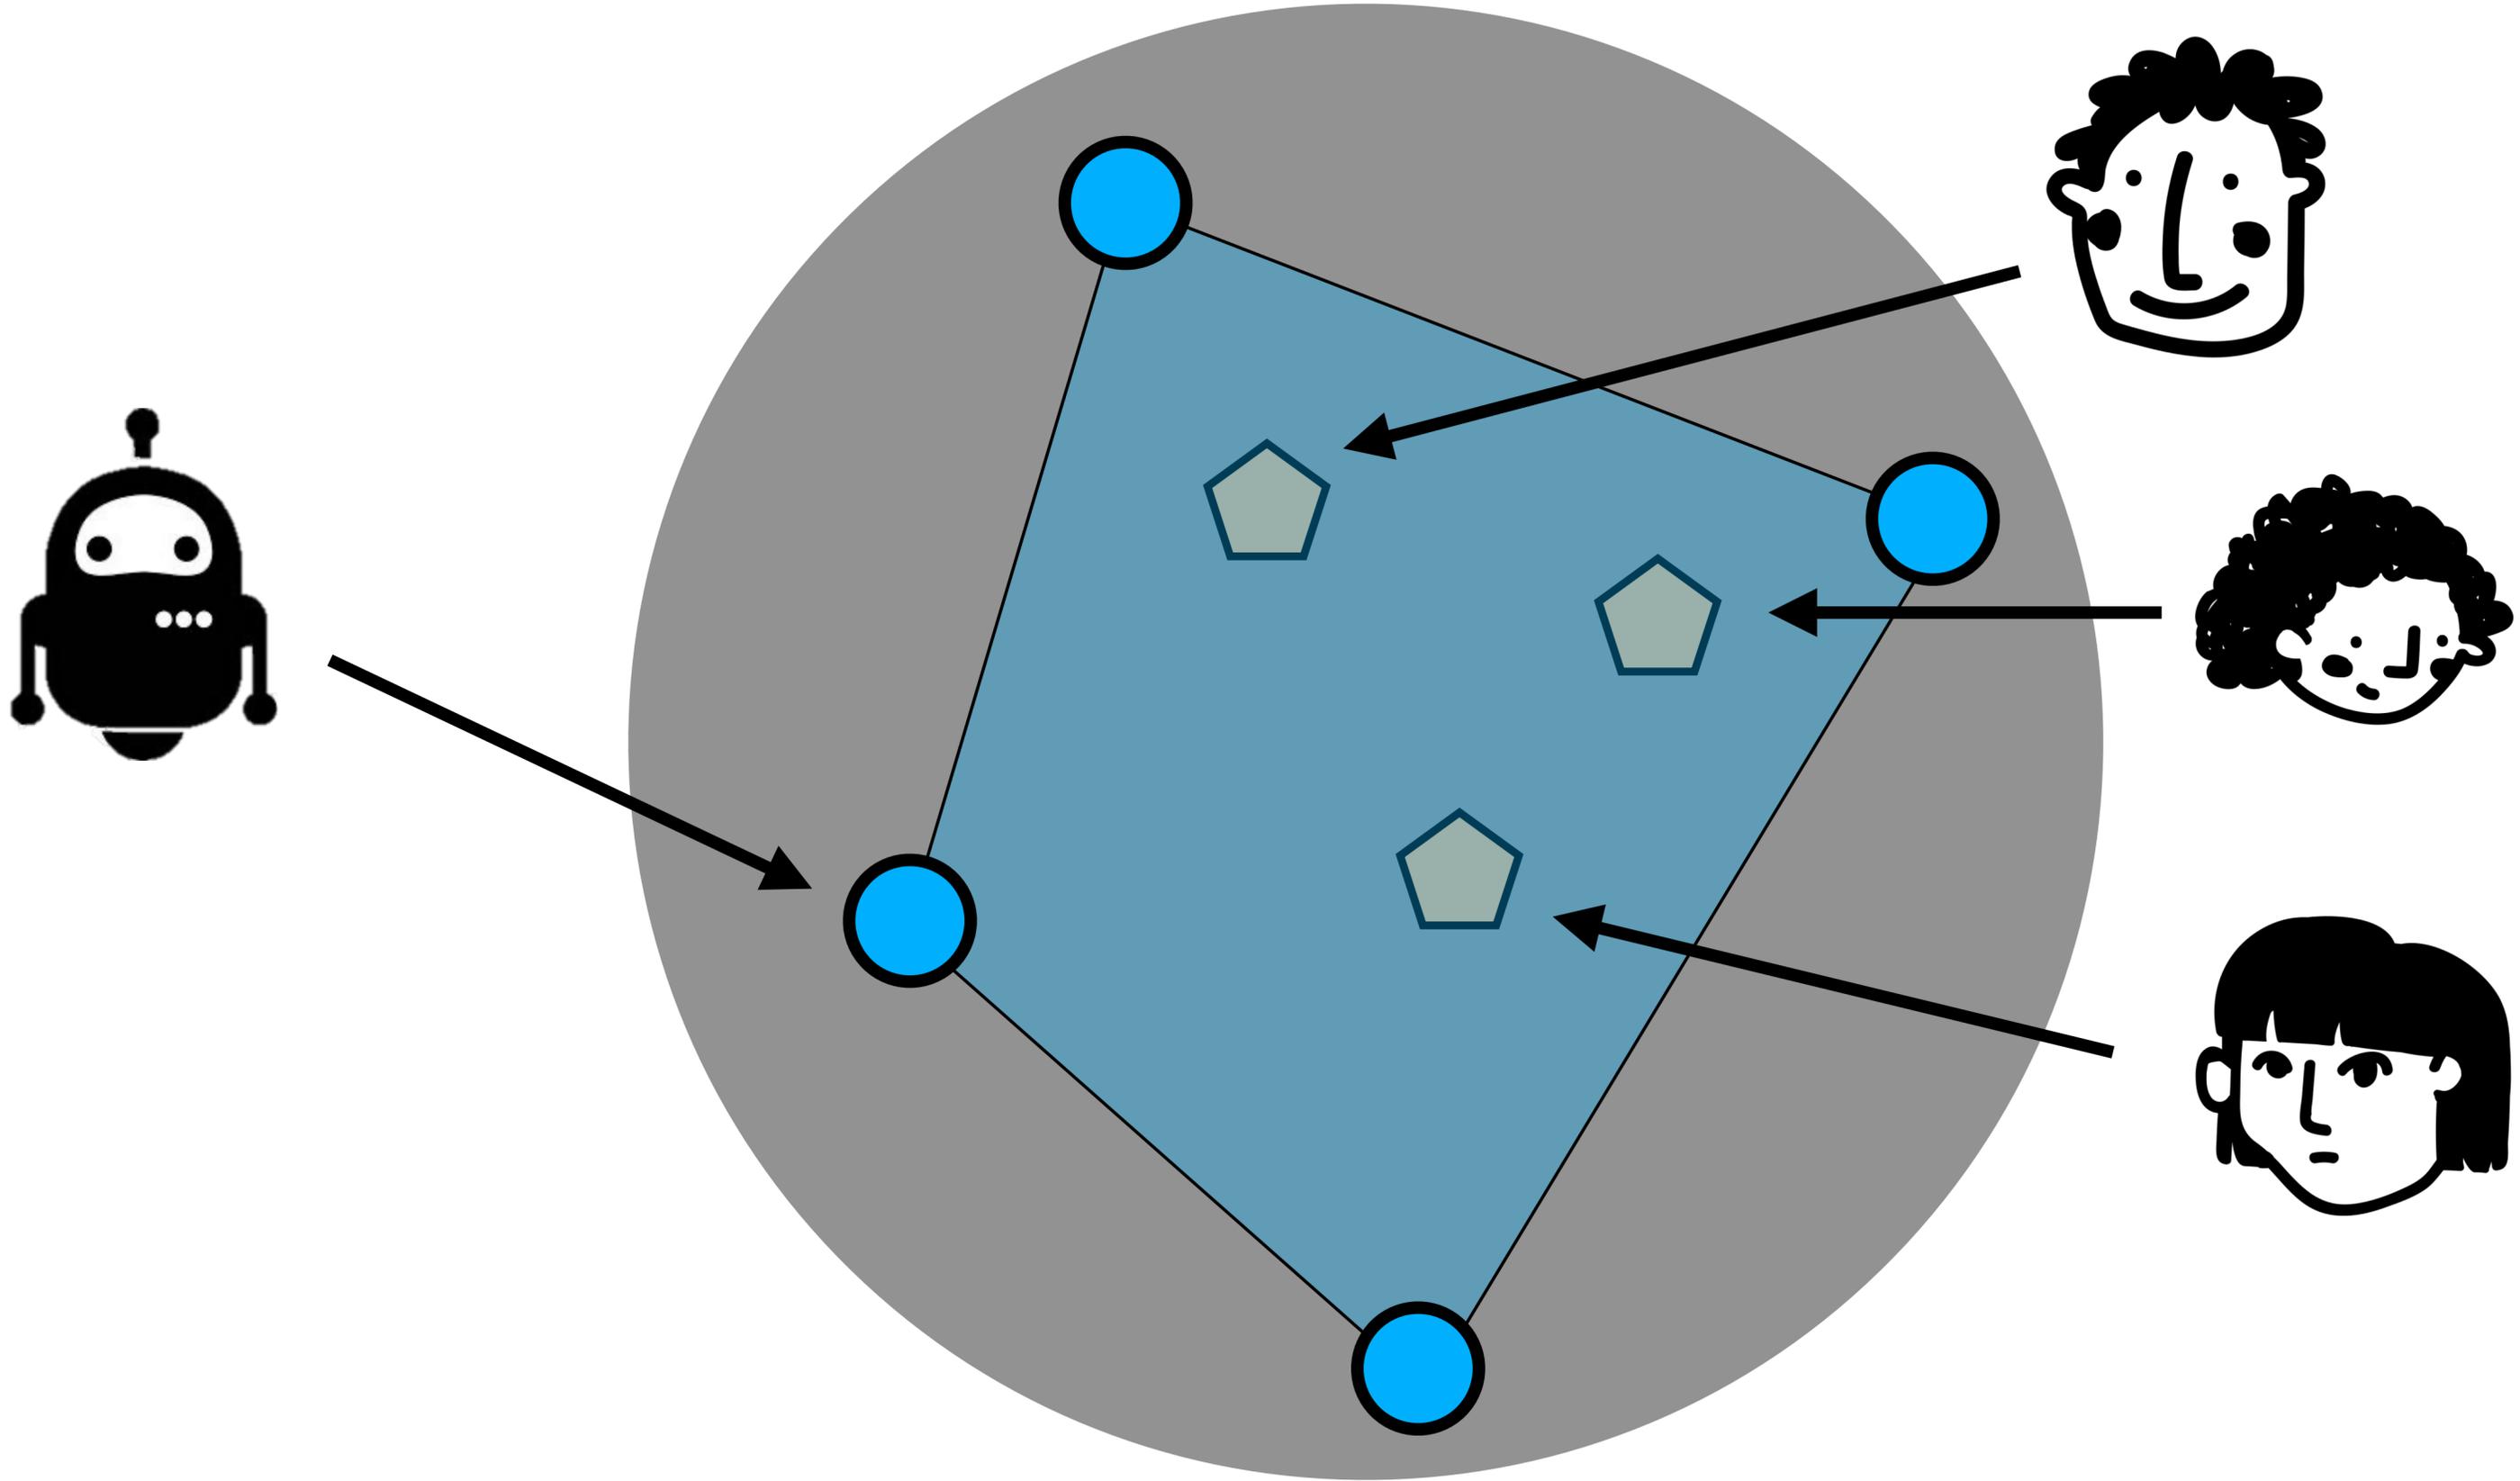
*Green is a human player.  
Blue is confused!*





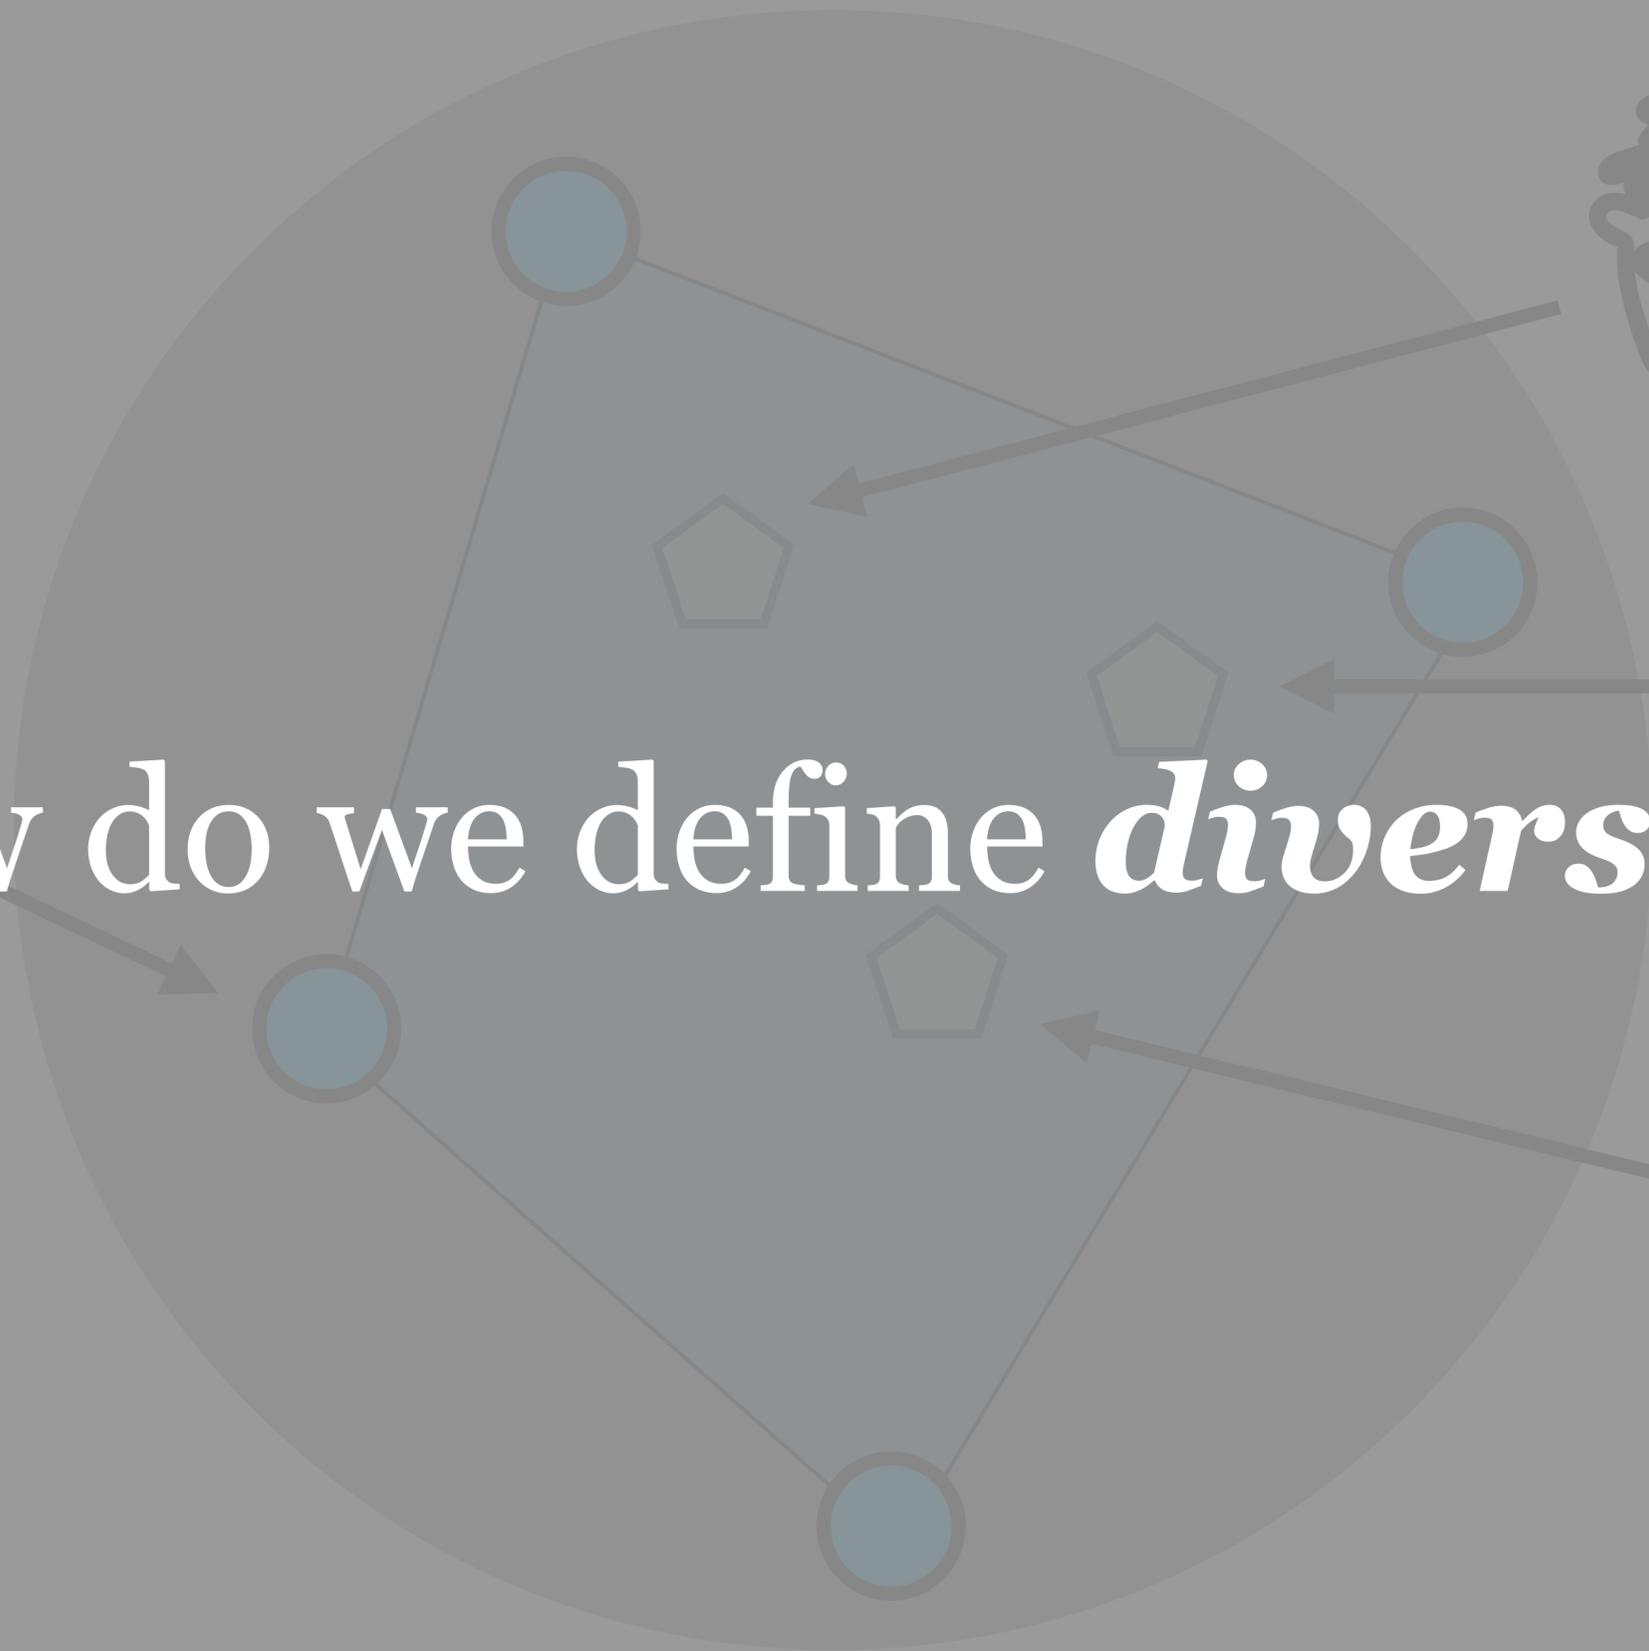








How do we define *diversity*?



# Statistical Diversity

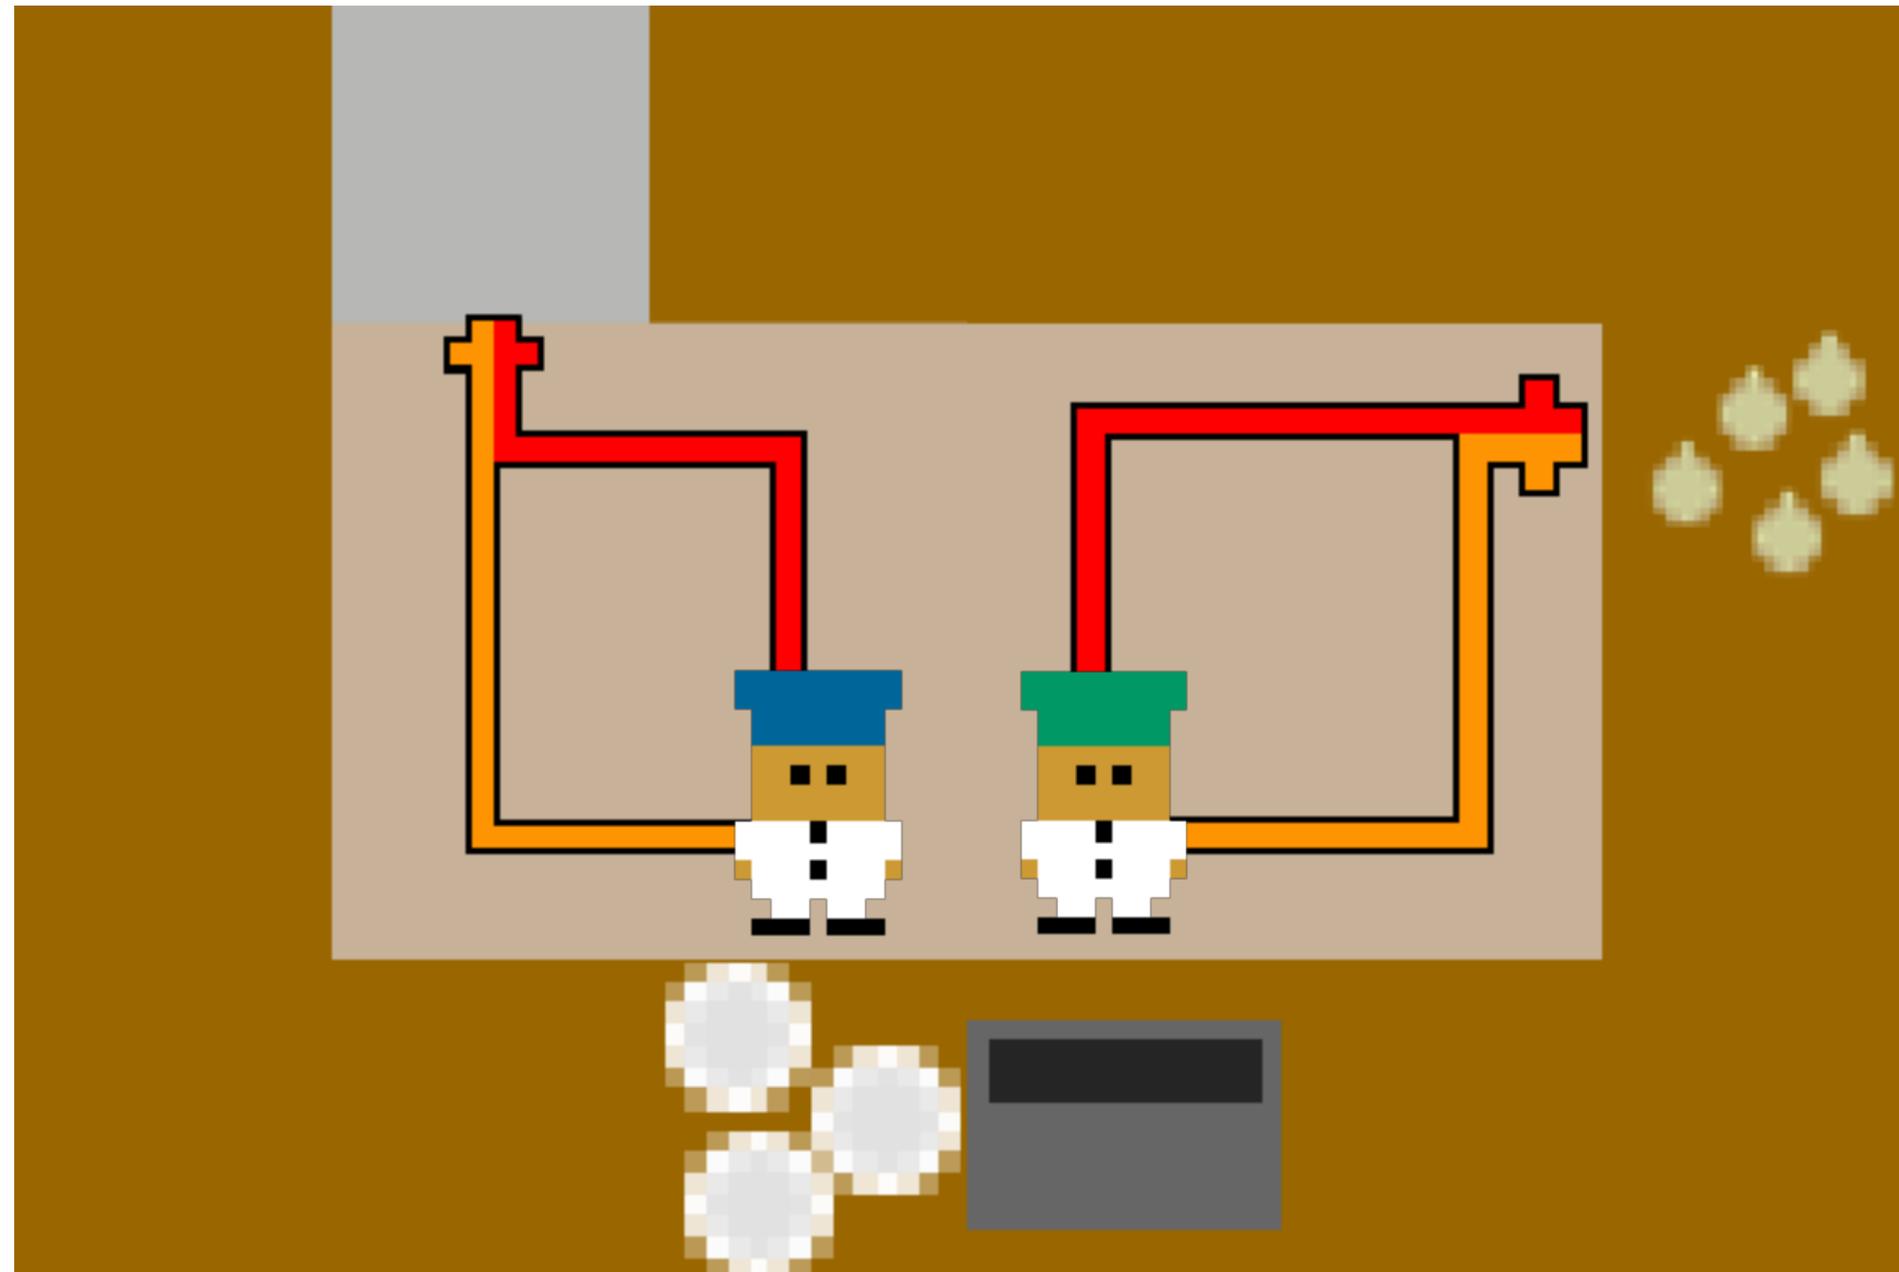
Maximize the difference in trajectories

# Statistical Diversity

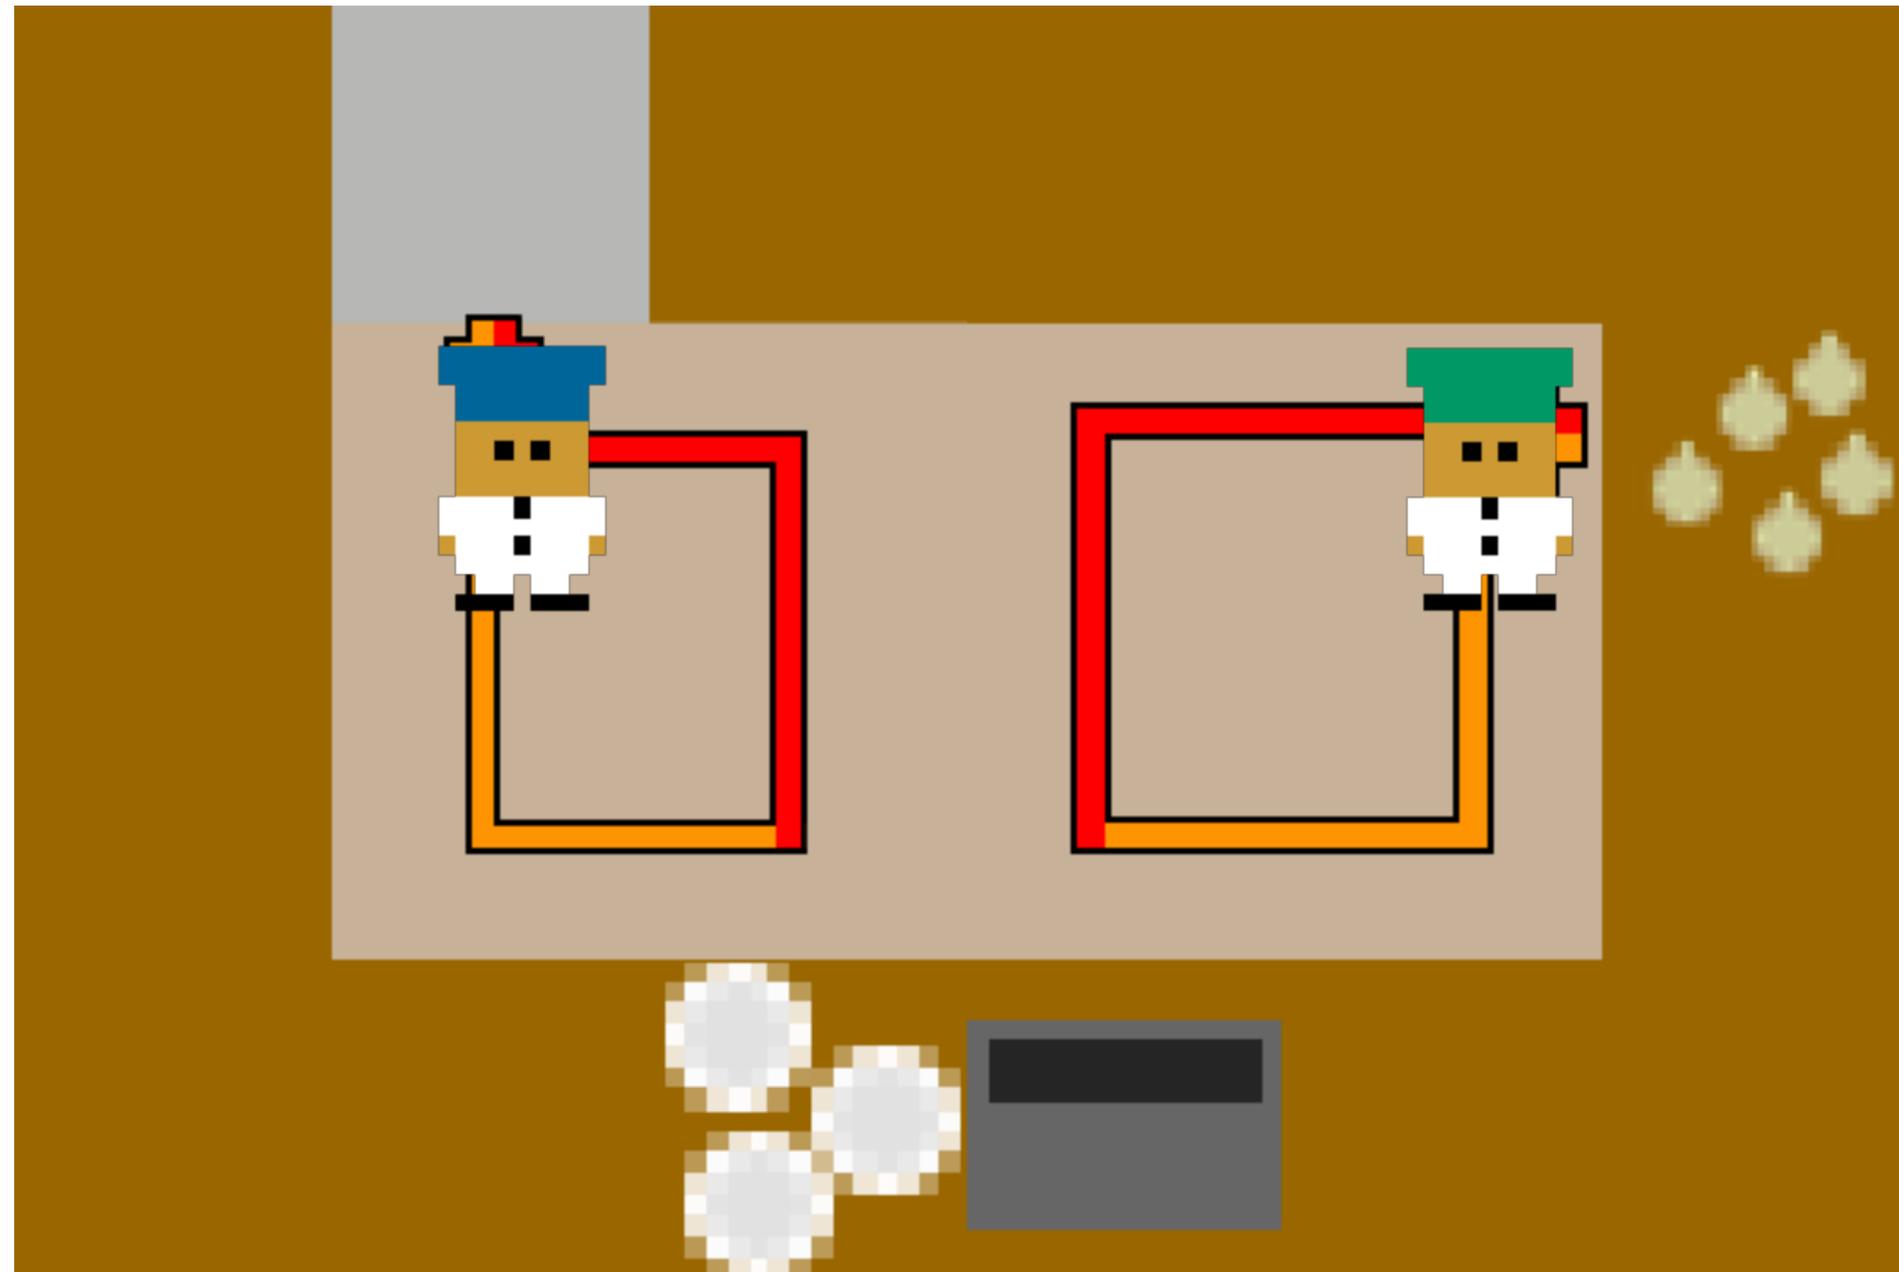
Maximize the difference in trajectories

$$\text{Maximize } C_{ADAP}(D) = \mathbb{E}_{s \in S} \left[ \mathbb{E}_{\substack{\pi_1, \pi_2 \in D \\ \pi_1 \neq \pi_2}} \exp(-D_{KL}(\pi_1(s), \pi_2(s))) \right]$$

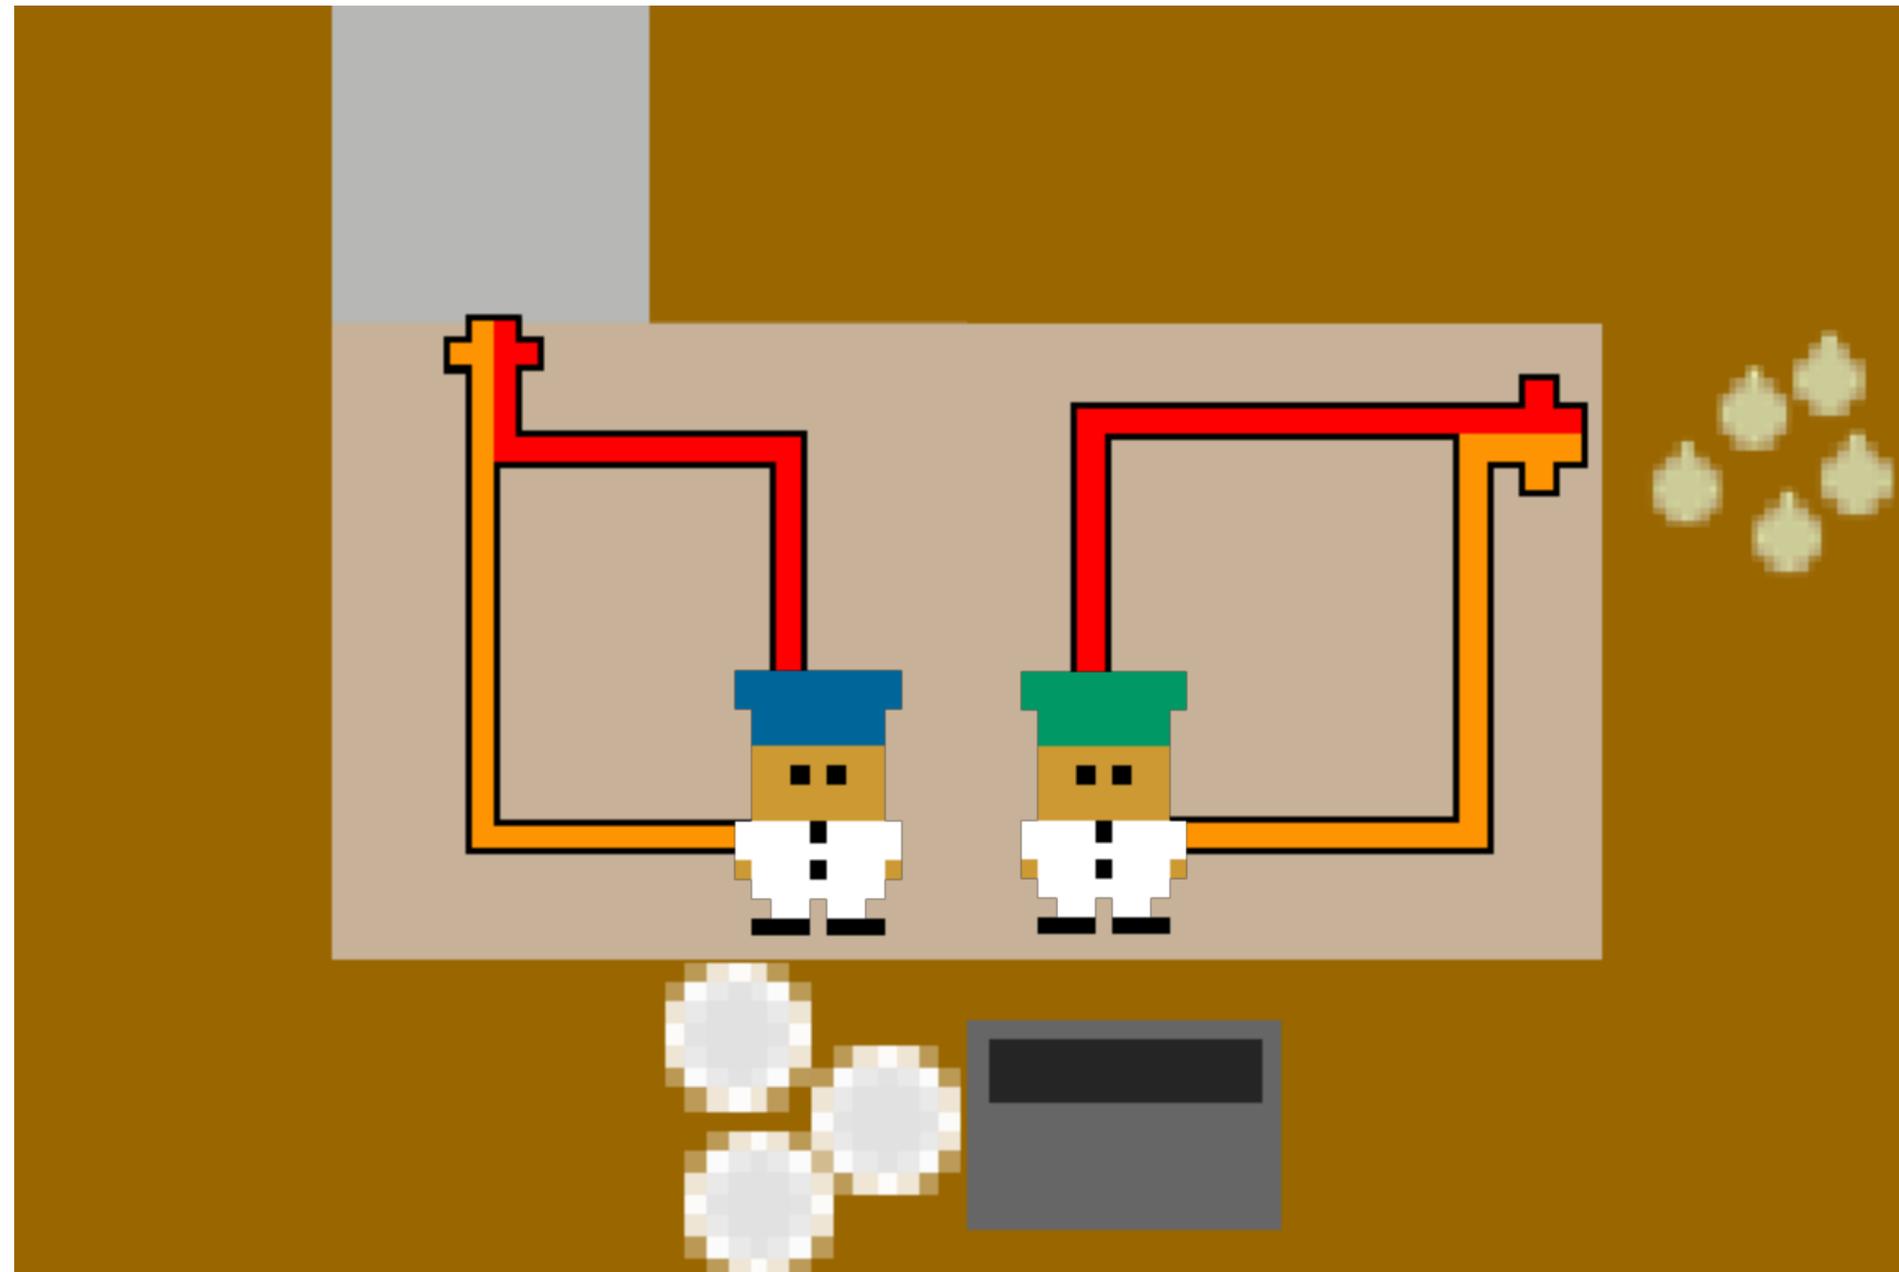
# Statistical Diversity



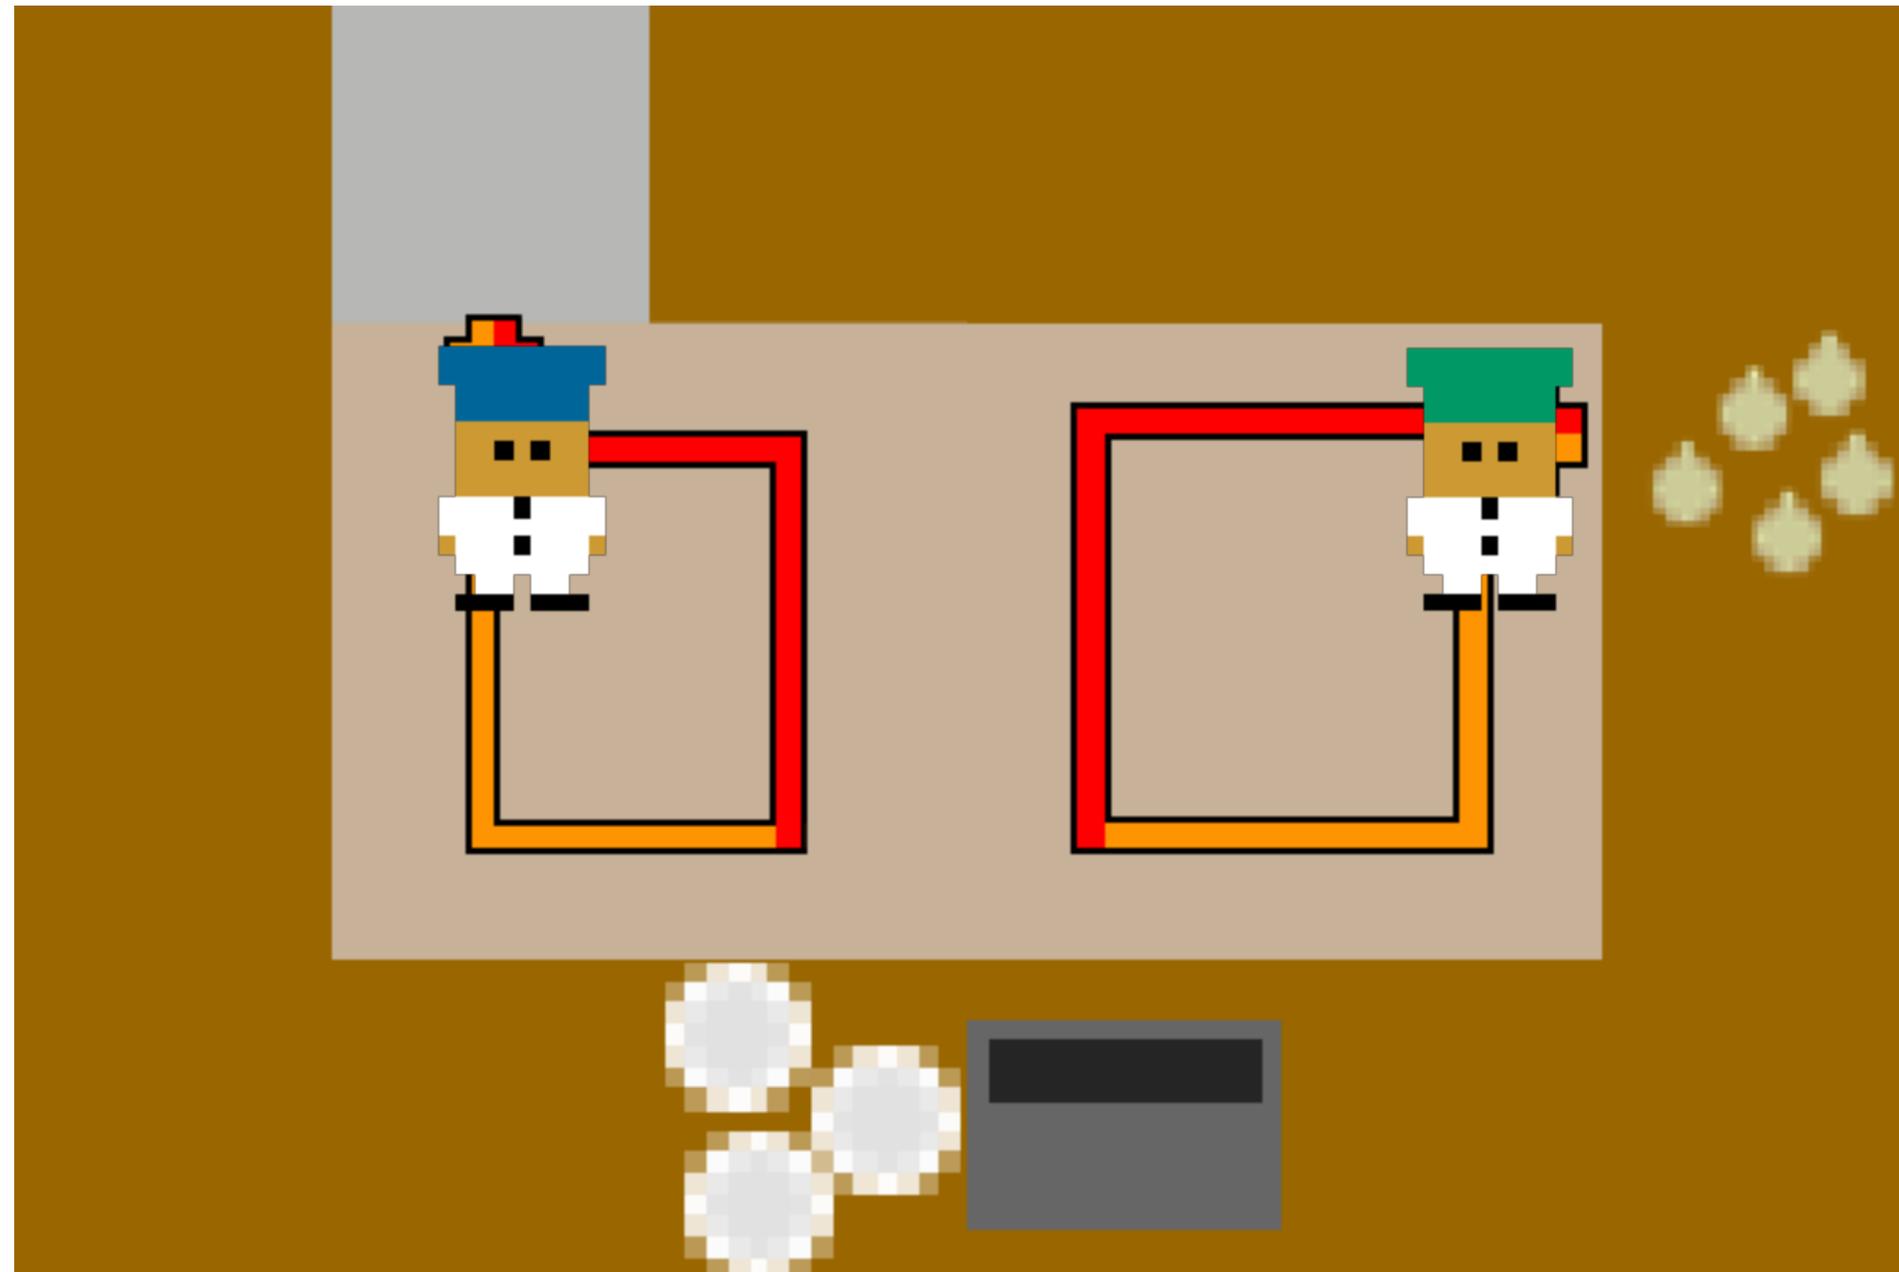
# Statistical Diversity



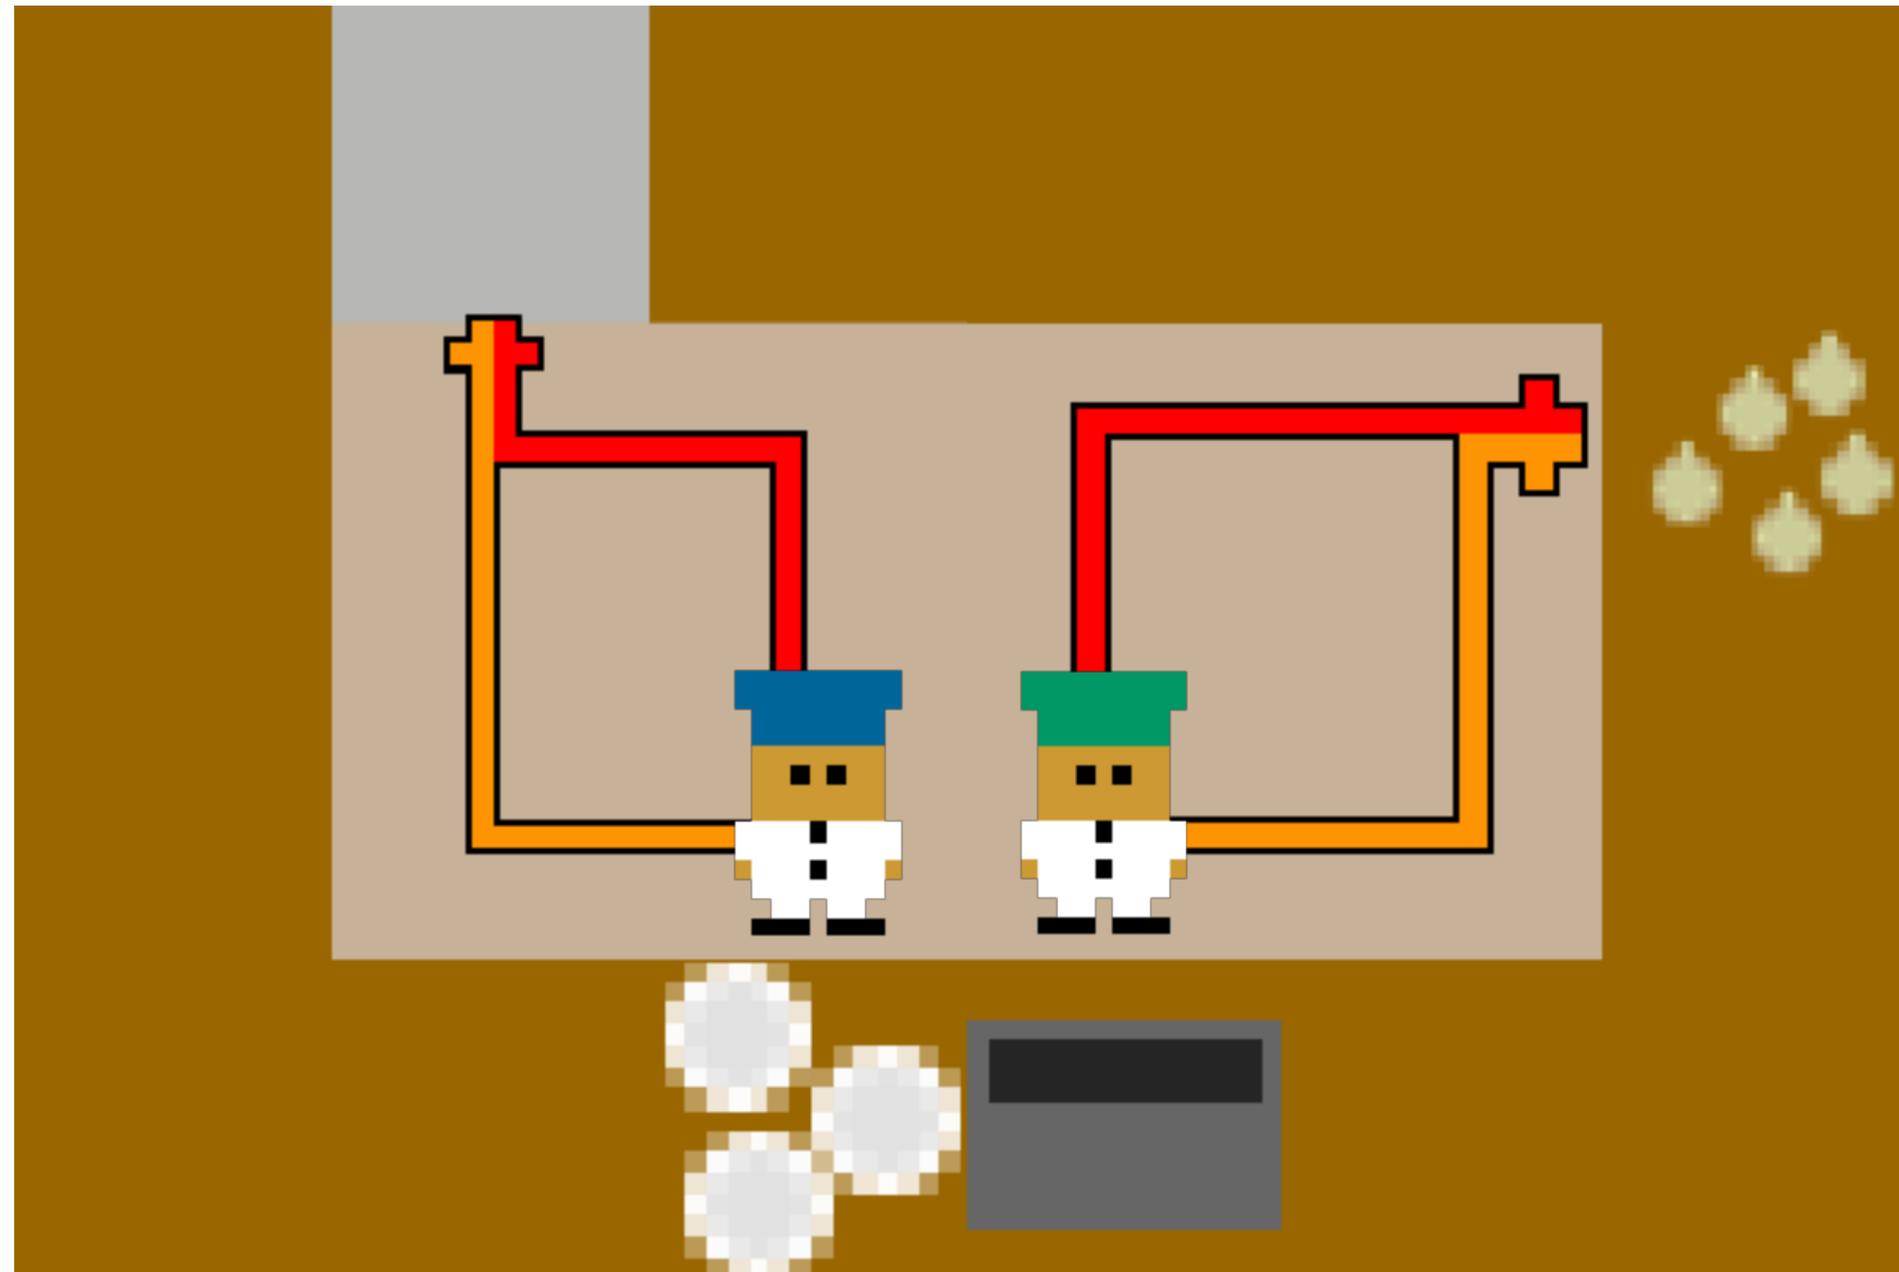
# Statistical Diversity



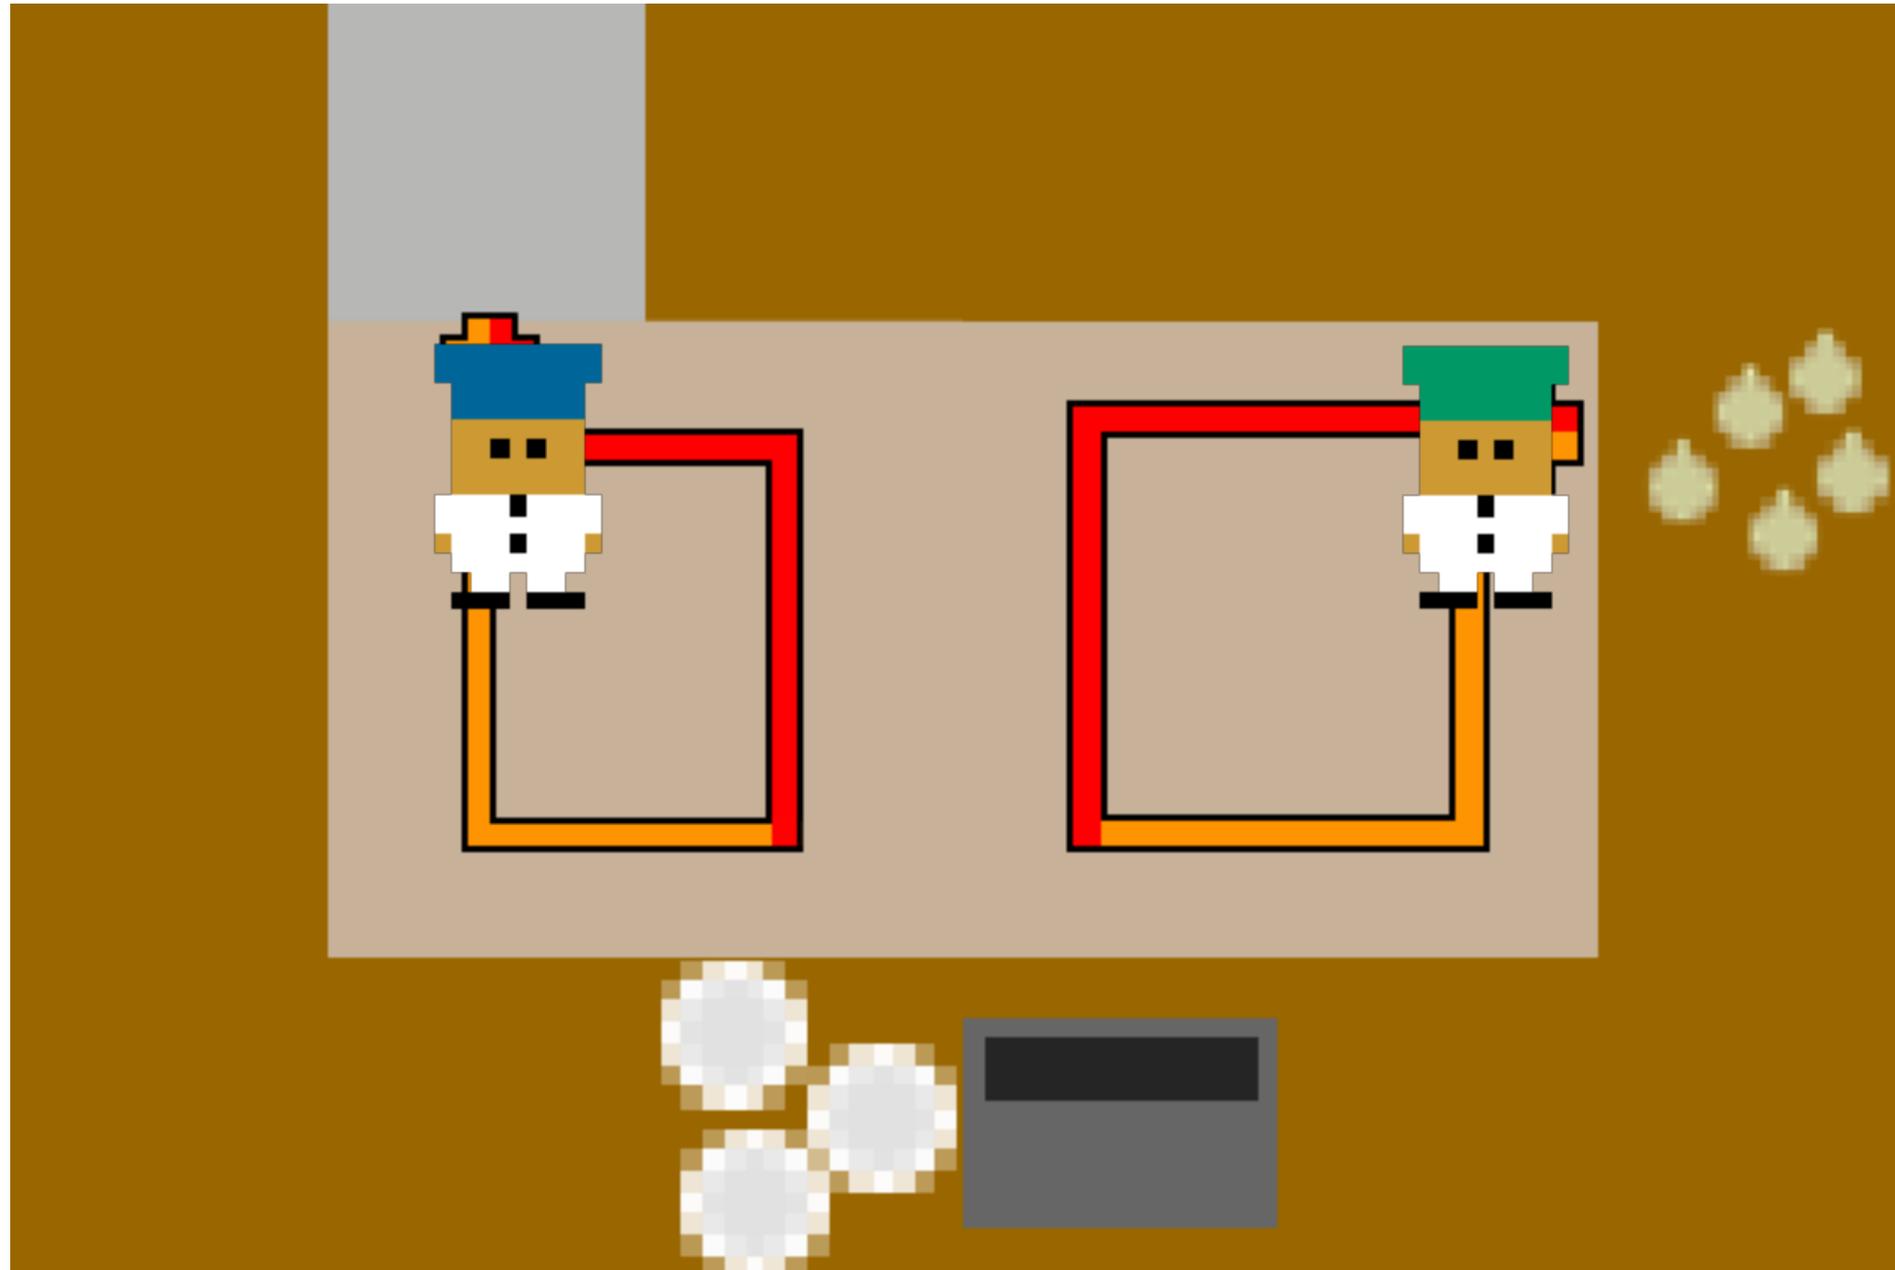
# Statistical Diversity



# Cross-Play

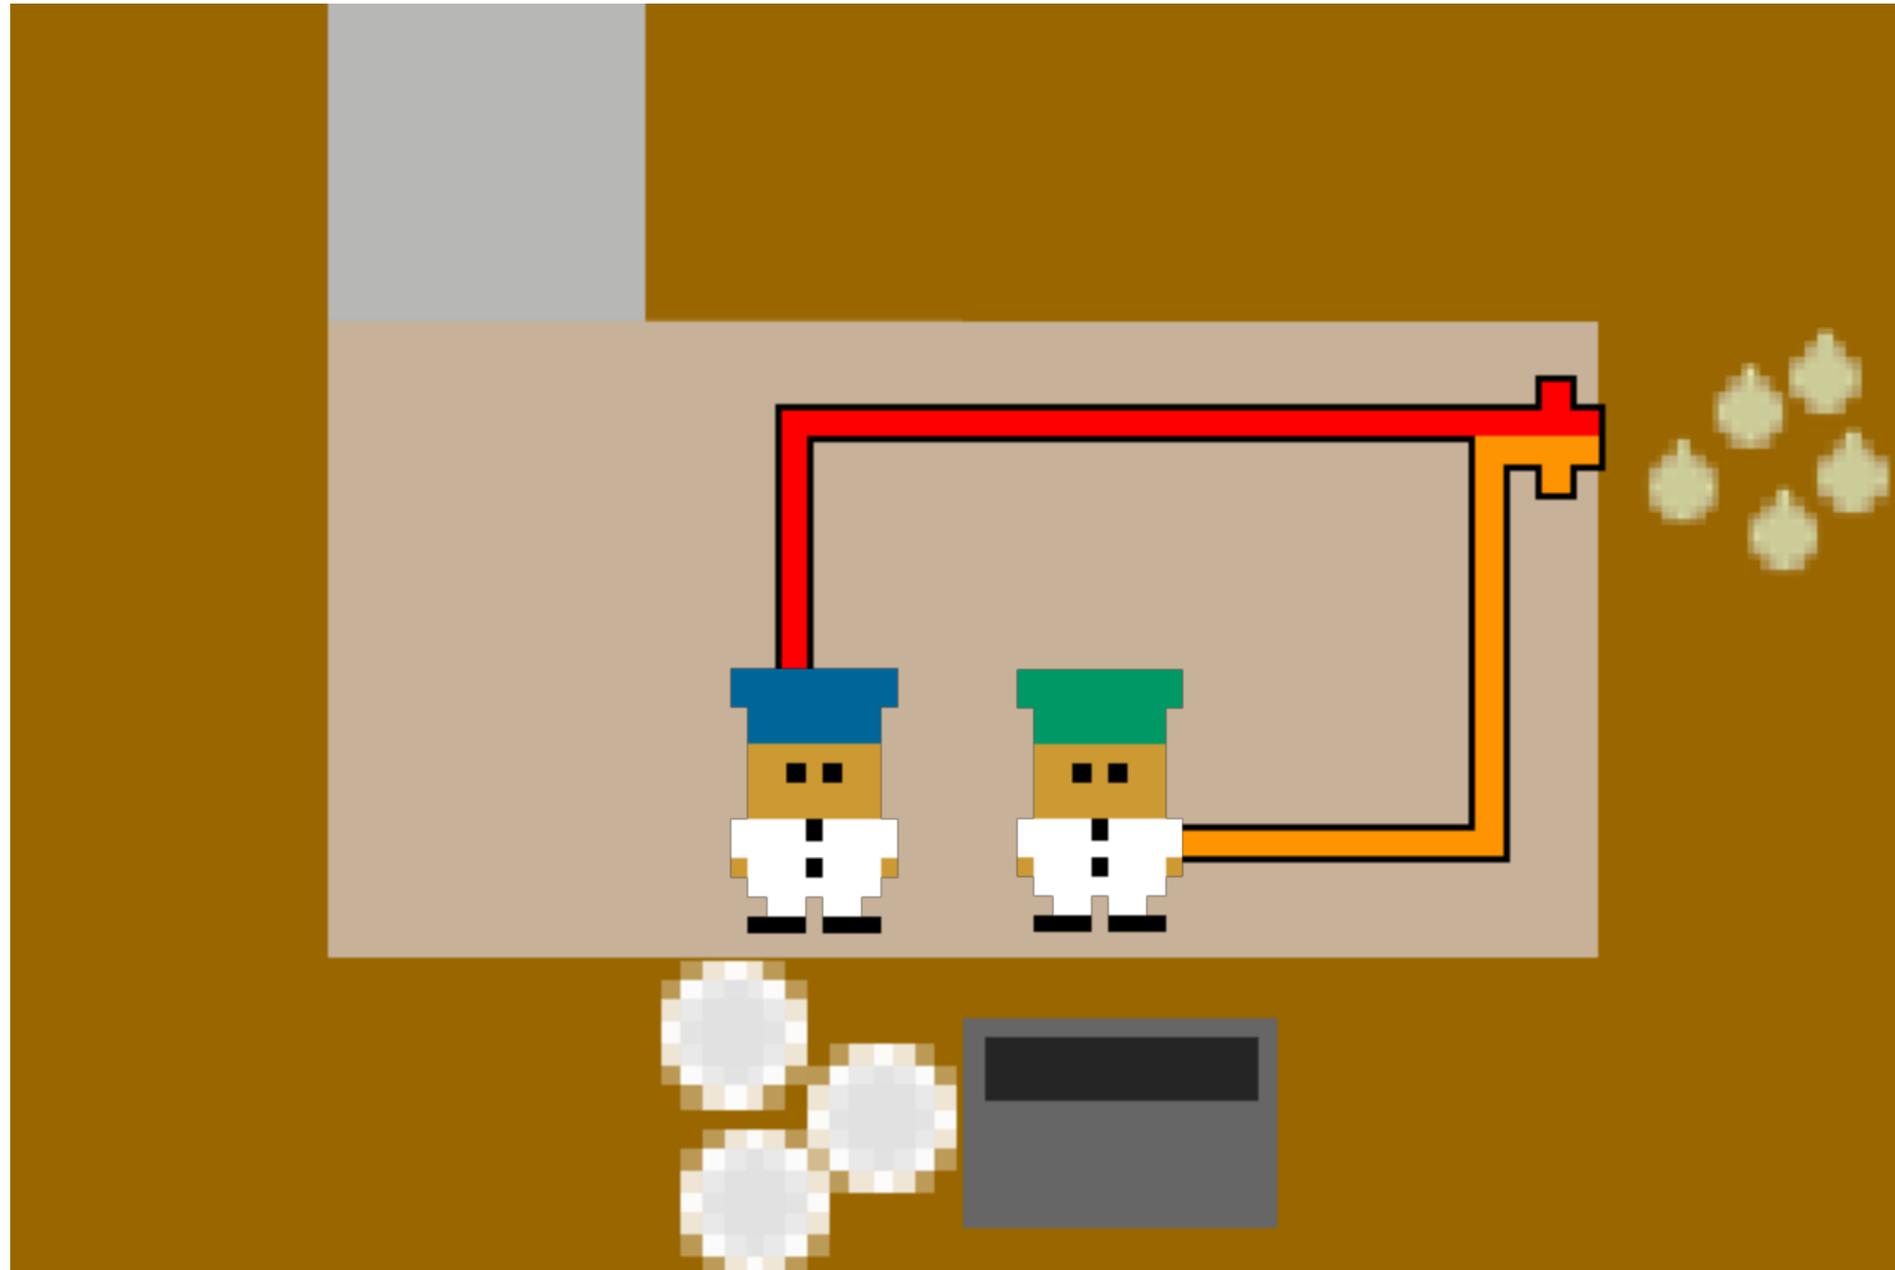


# Cross-Play

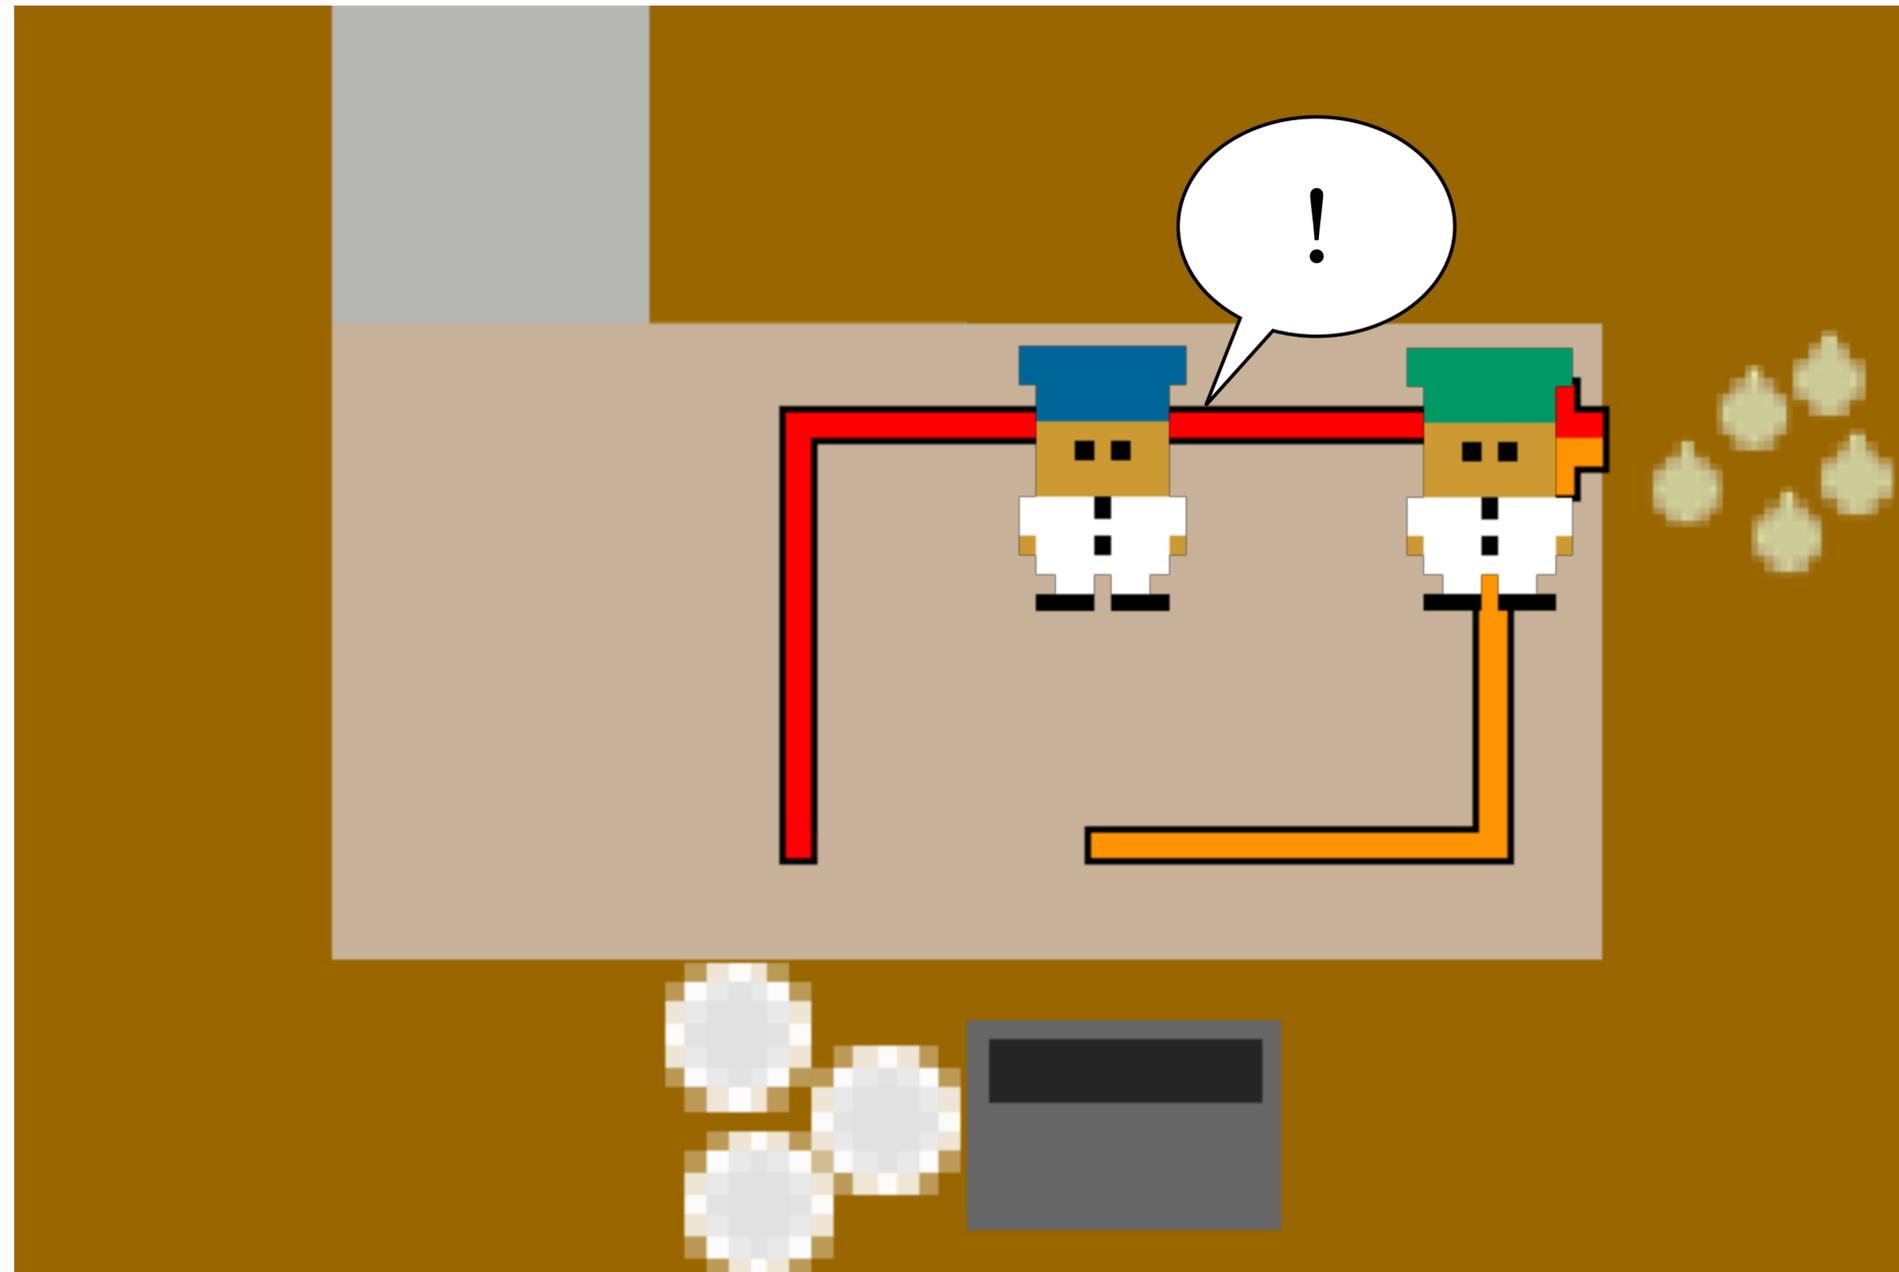


These policies are *compatible*, and are therefore *similar*

# Cross-Play



# Cross-Play



These policies are *incompatible*, and are therefore *different*

# Cross-Play

Minimize the cross-play score between different conventions

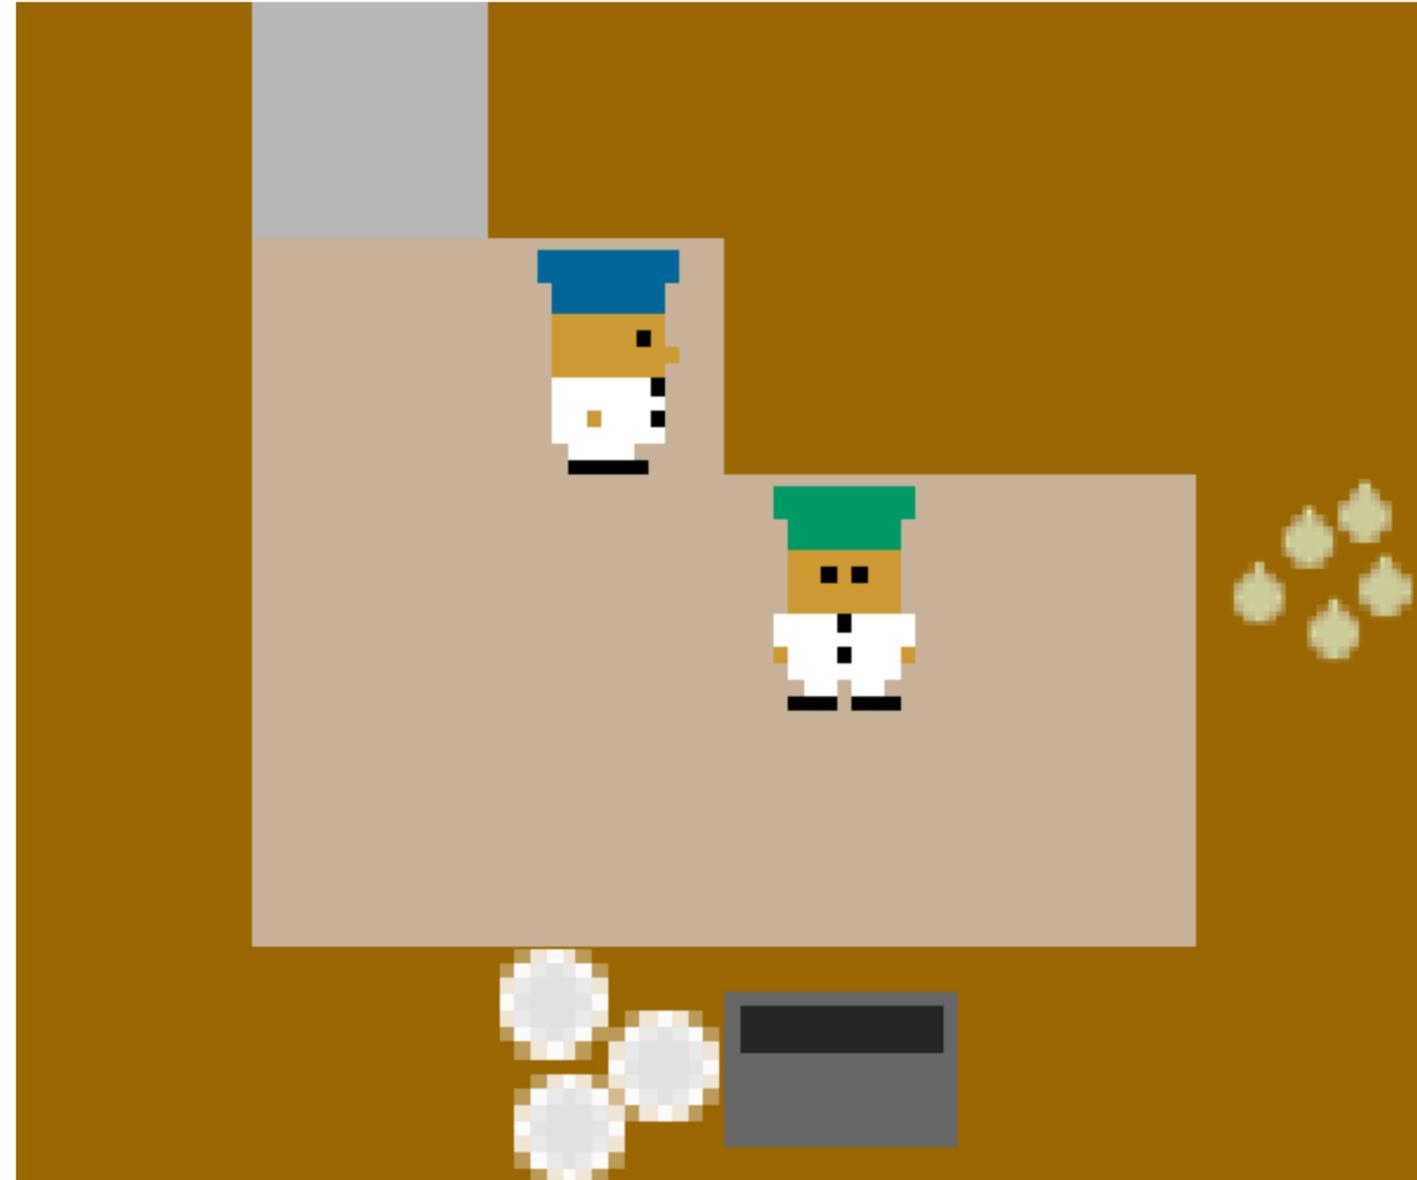
# Cross-Play

Minimize the cross-play score between different conventions

$$\textit{Minimize} \quad L(\pi_2) = -J(\pi_2, \pi_2) + J(\pi_2, \pi_1)$$

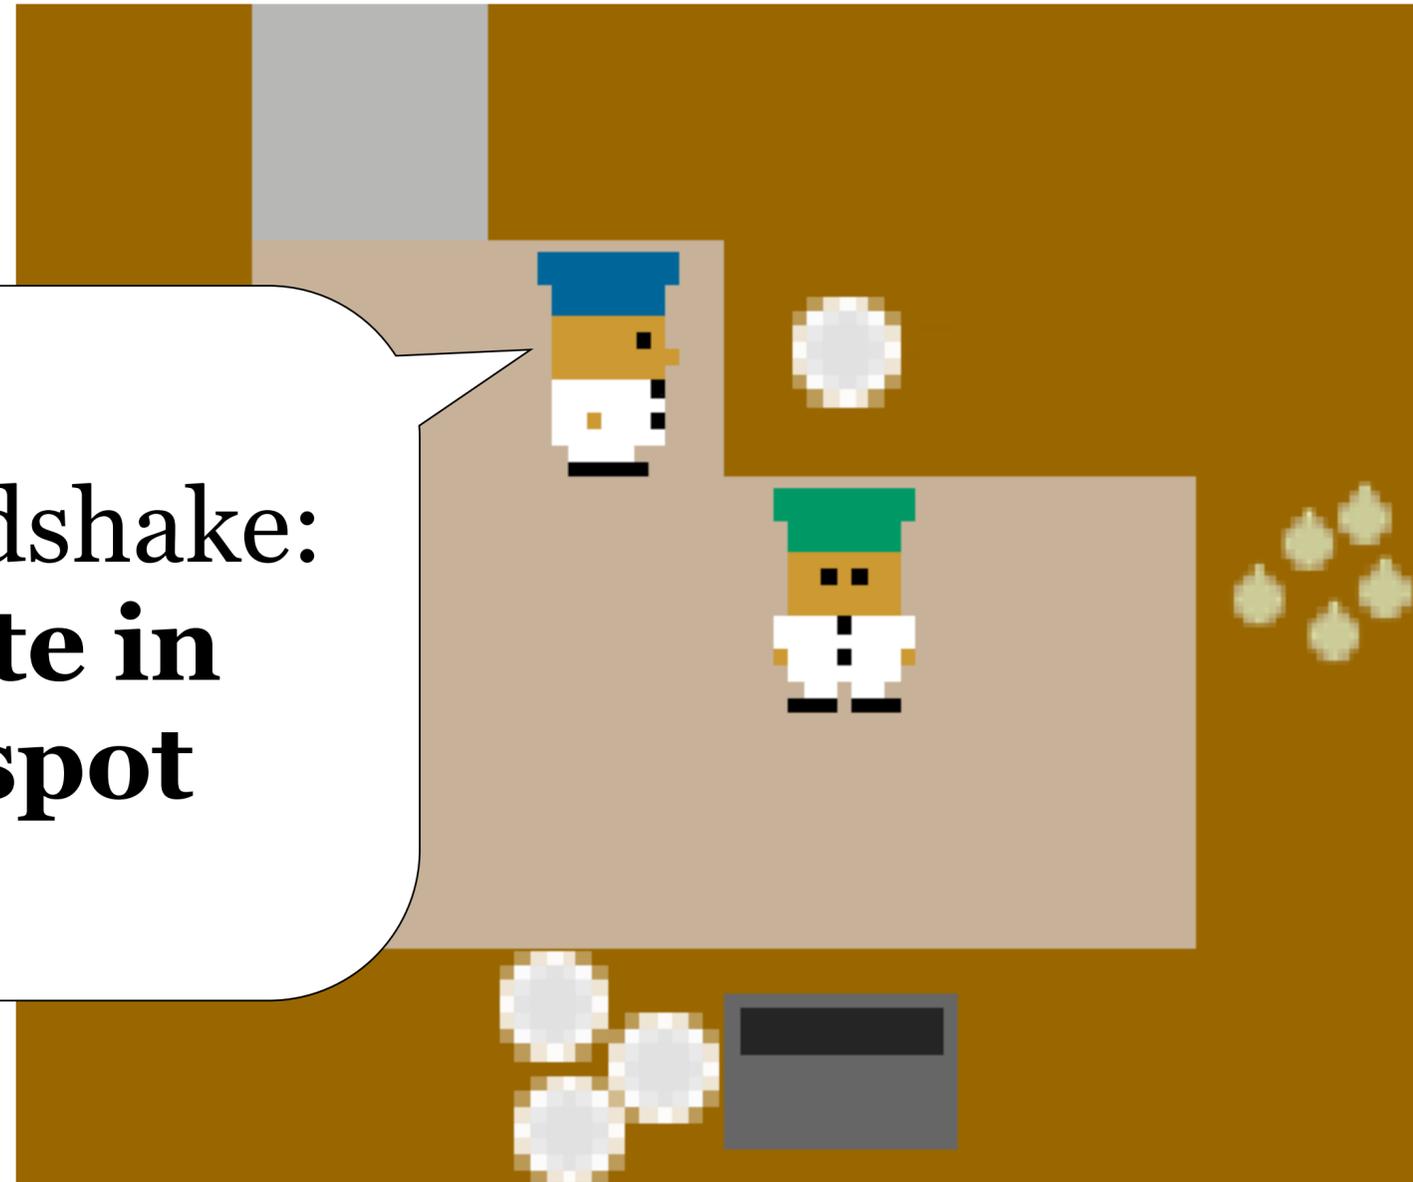
Where  $J$  is the expected reward when two policies play together.

# Handshakes

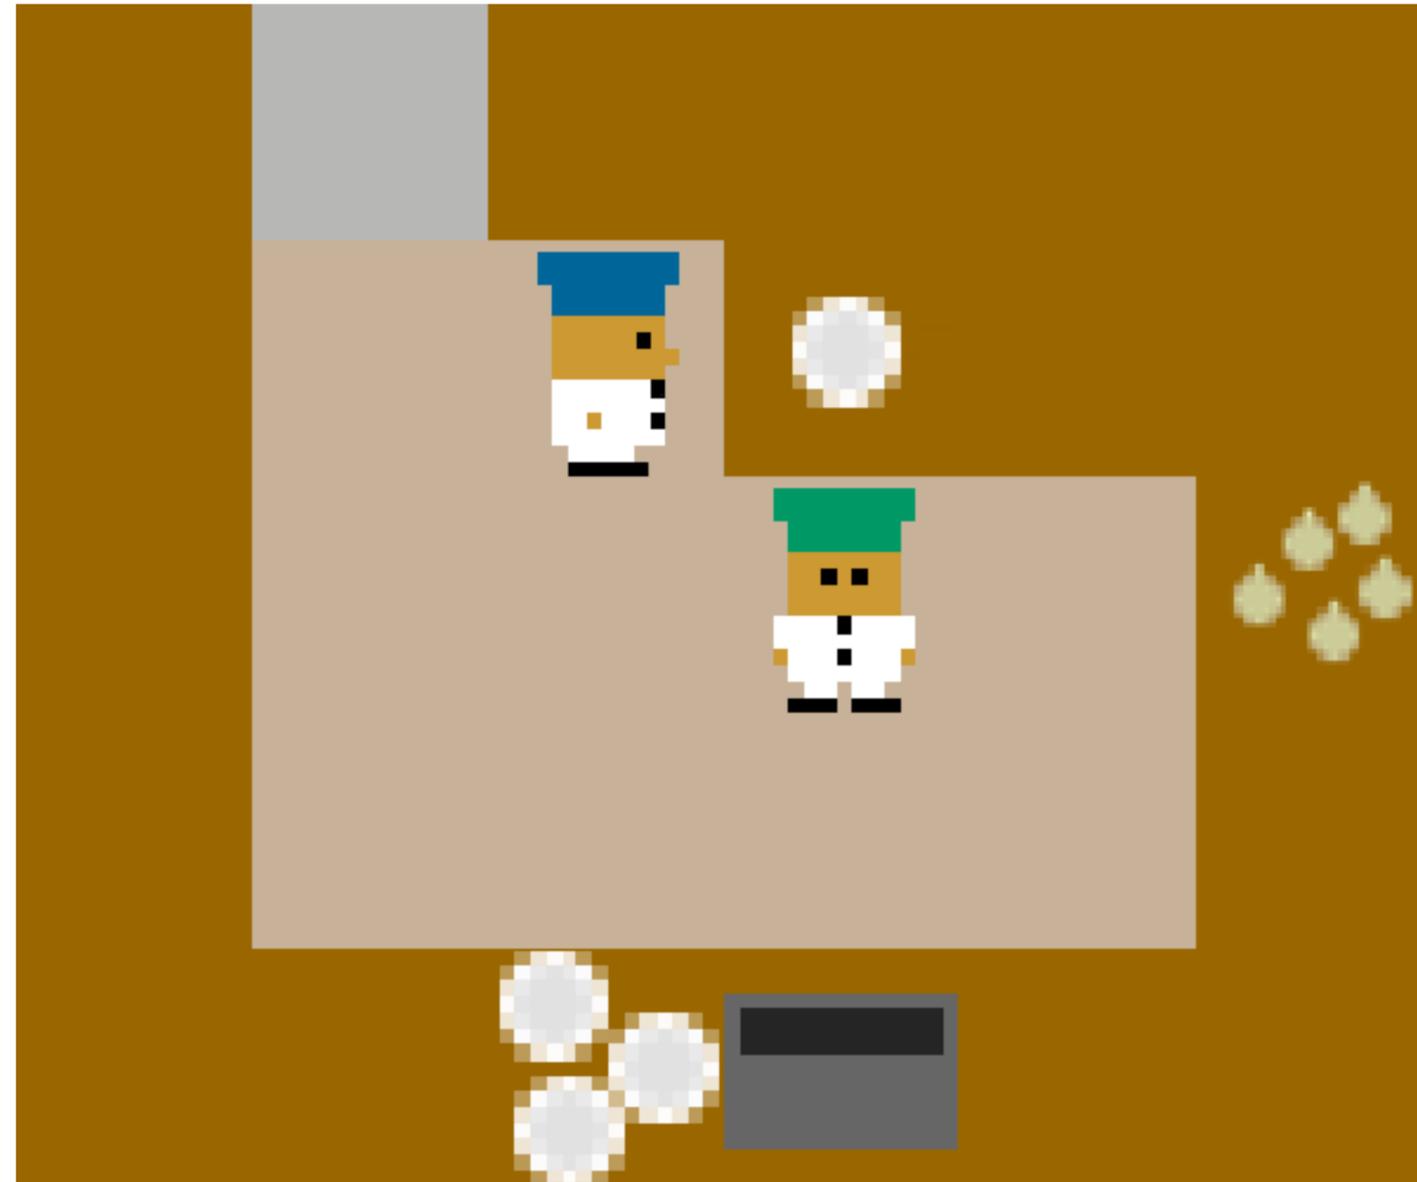


# Handshakes

Initiate Handshake:  
**Place plate in  
strange spot**



# Handshakes

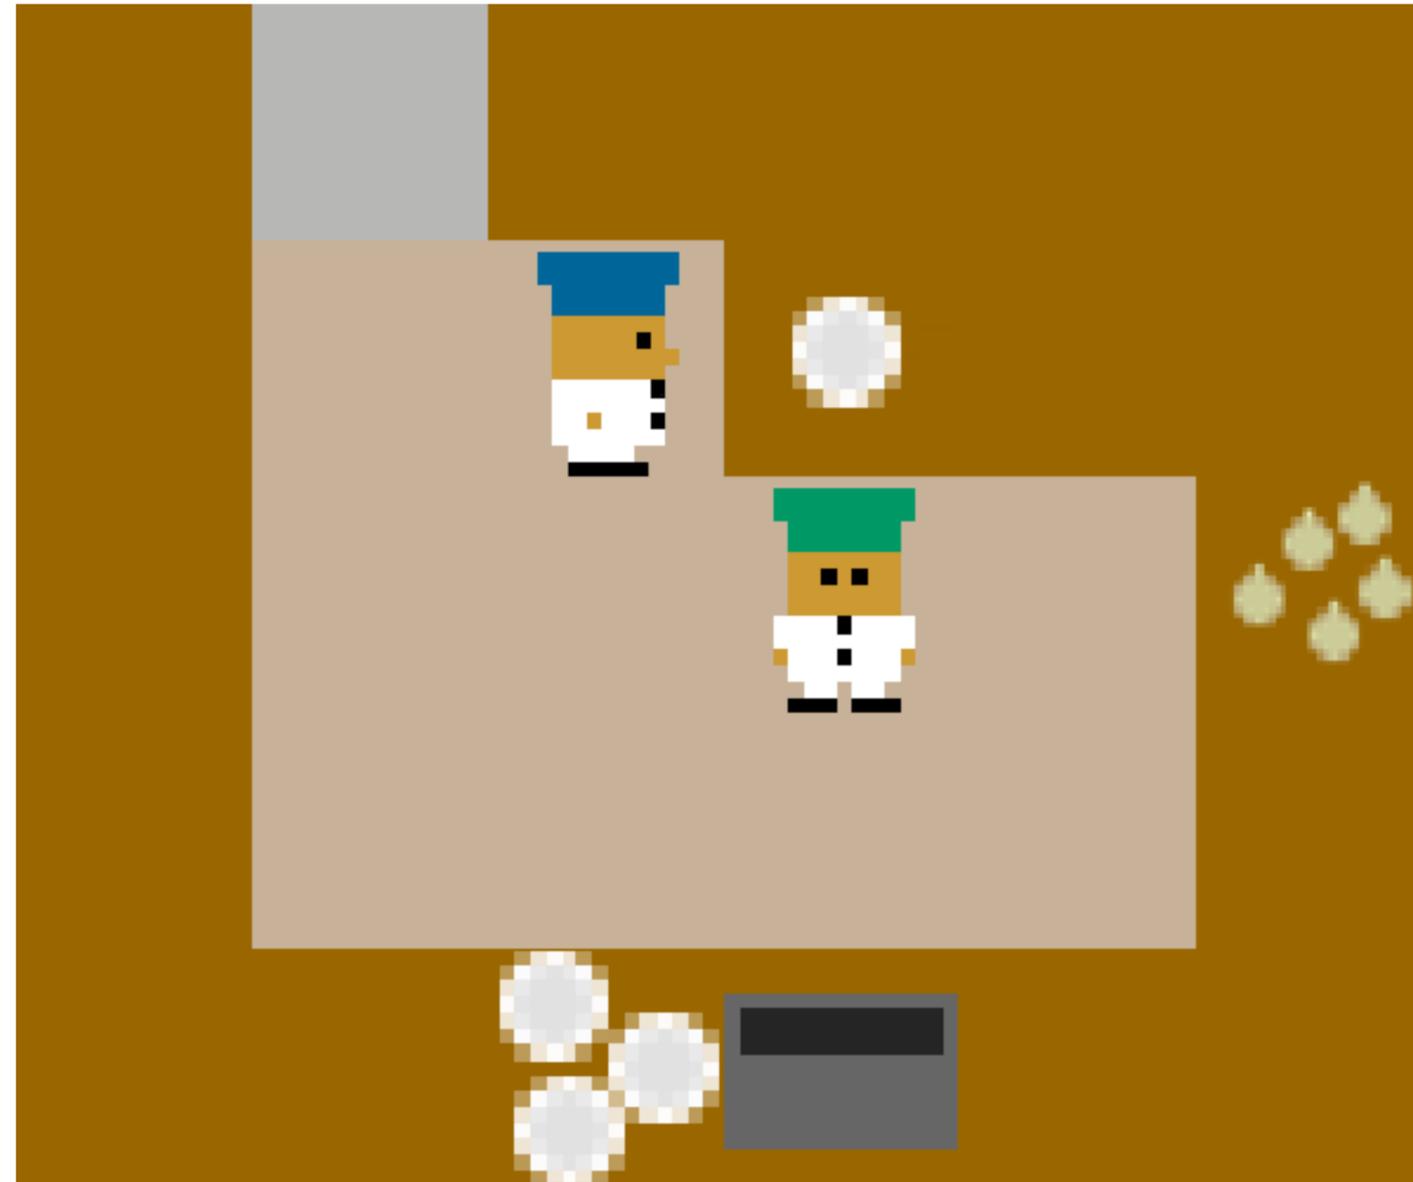


# Handshakes

Handshake  
acknowledged!  
**I'll work hard!**



# Handshakes

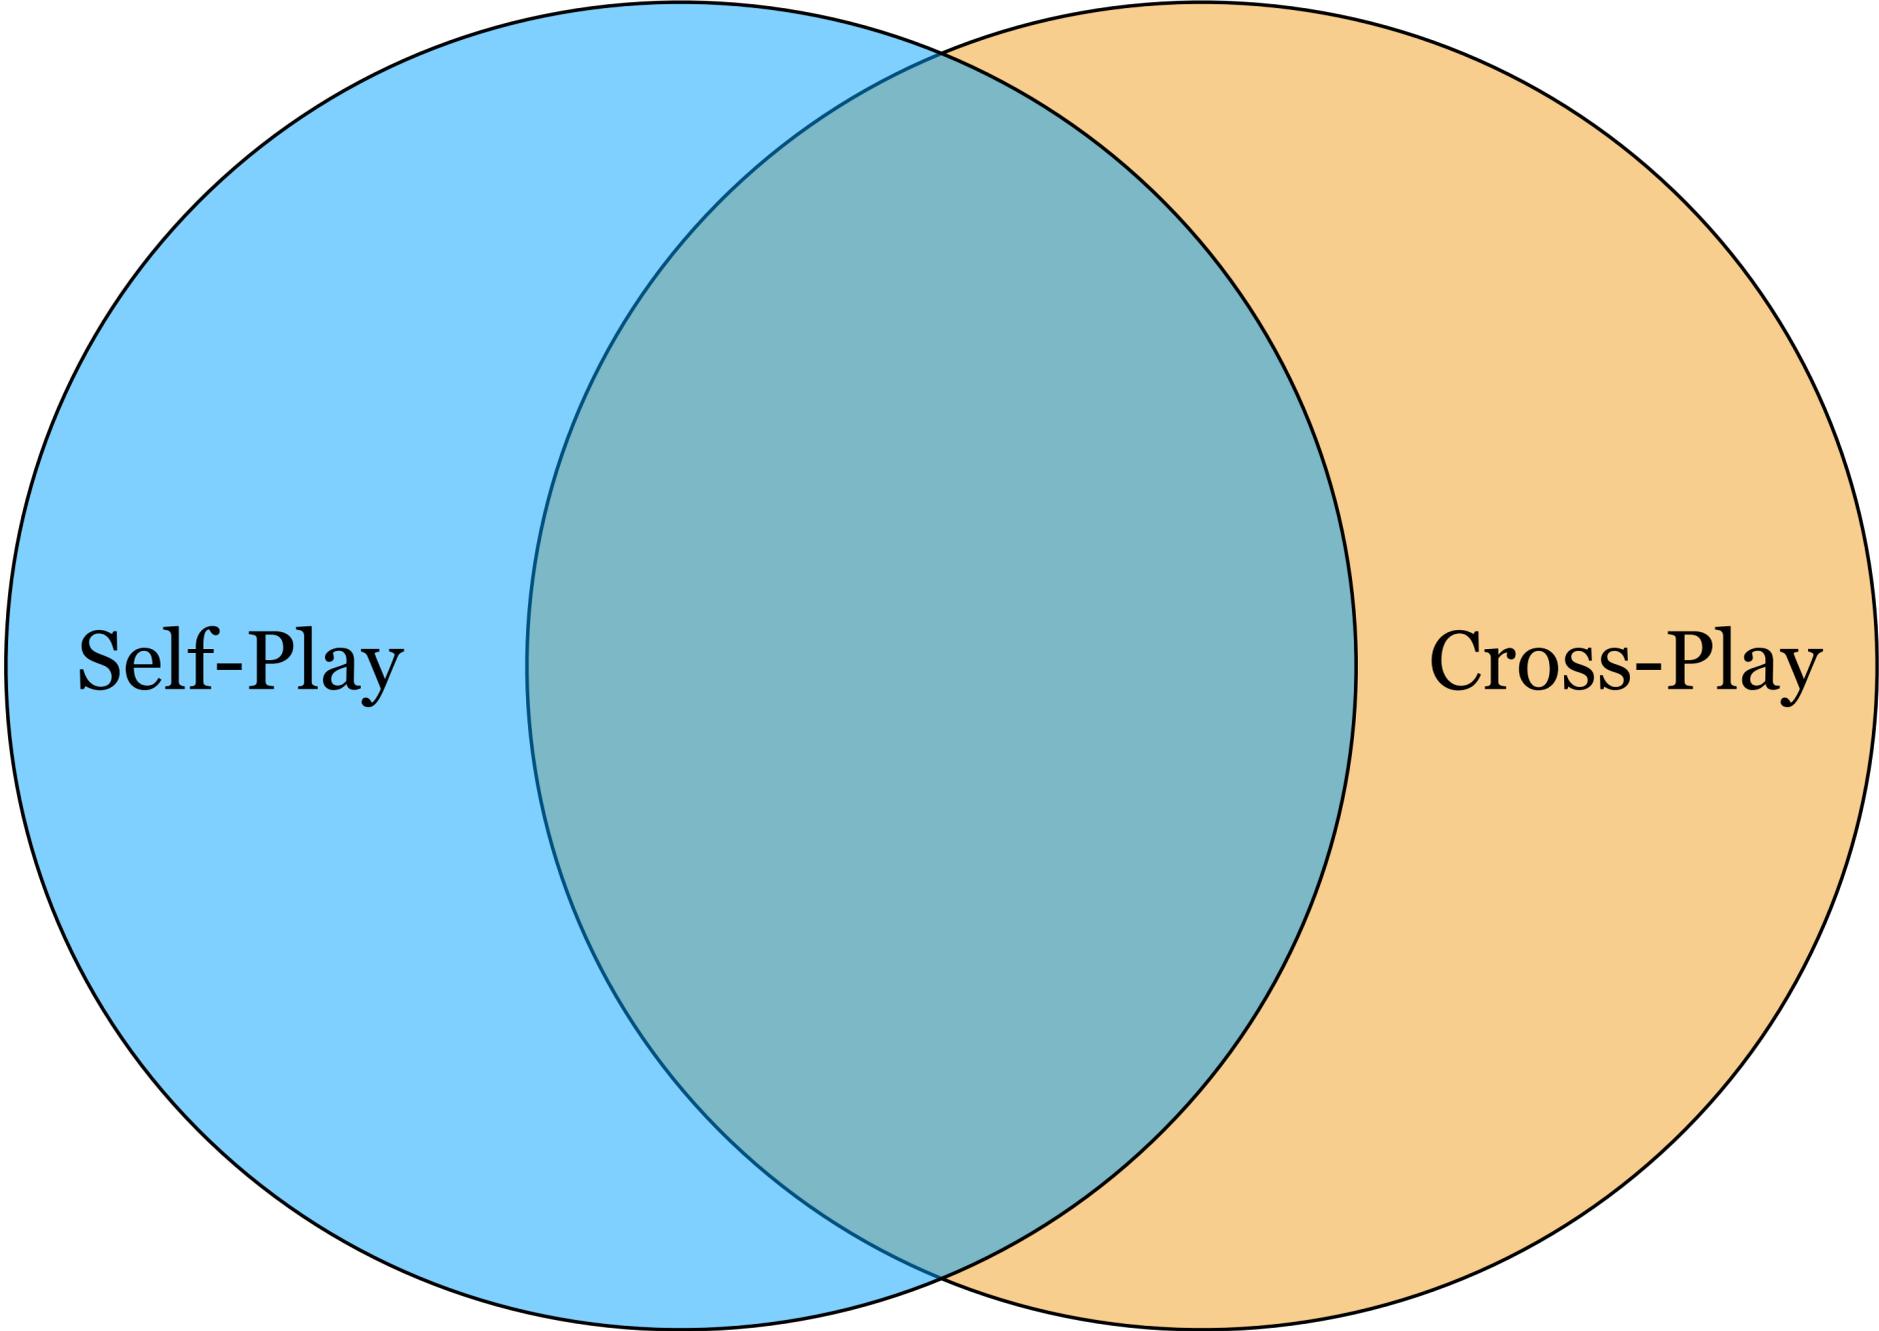


# Handshakes

Handshake rejected!  
**I won't help.**

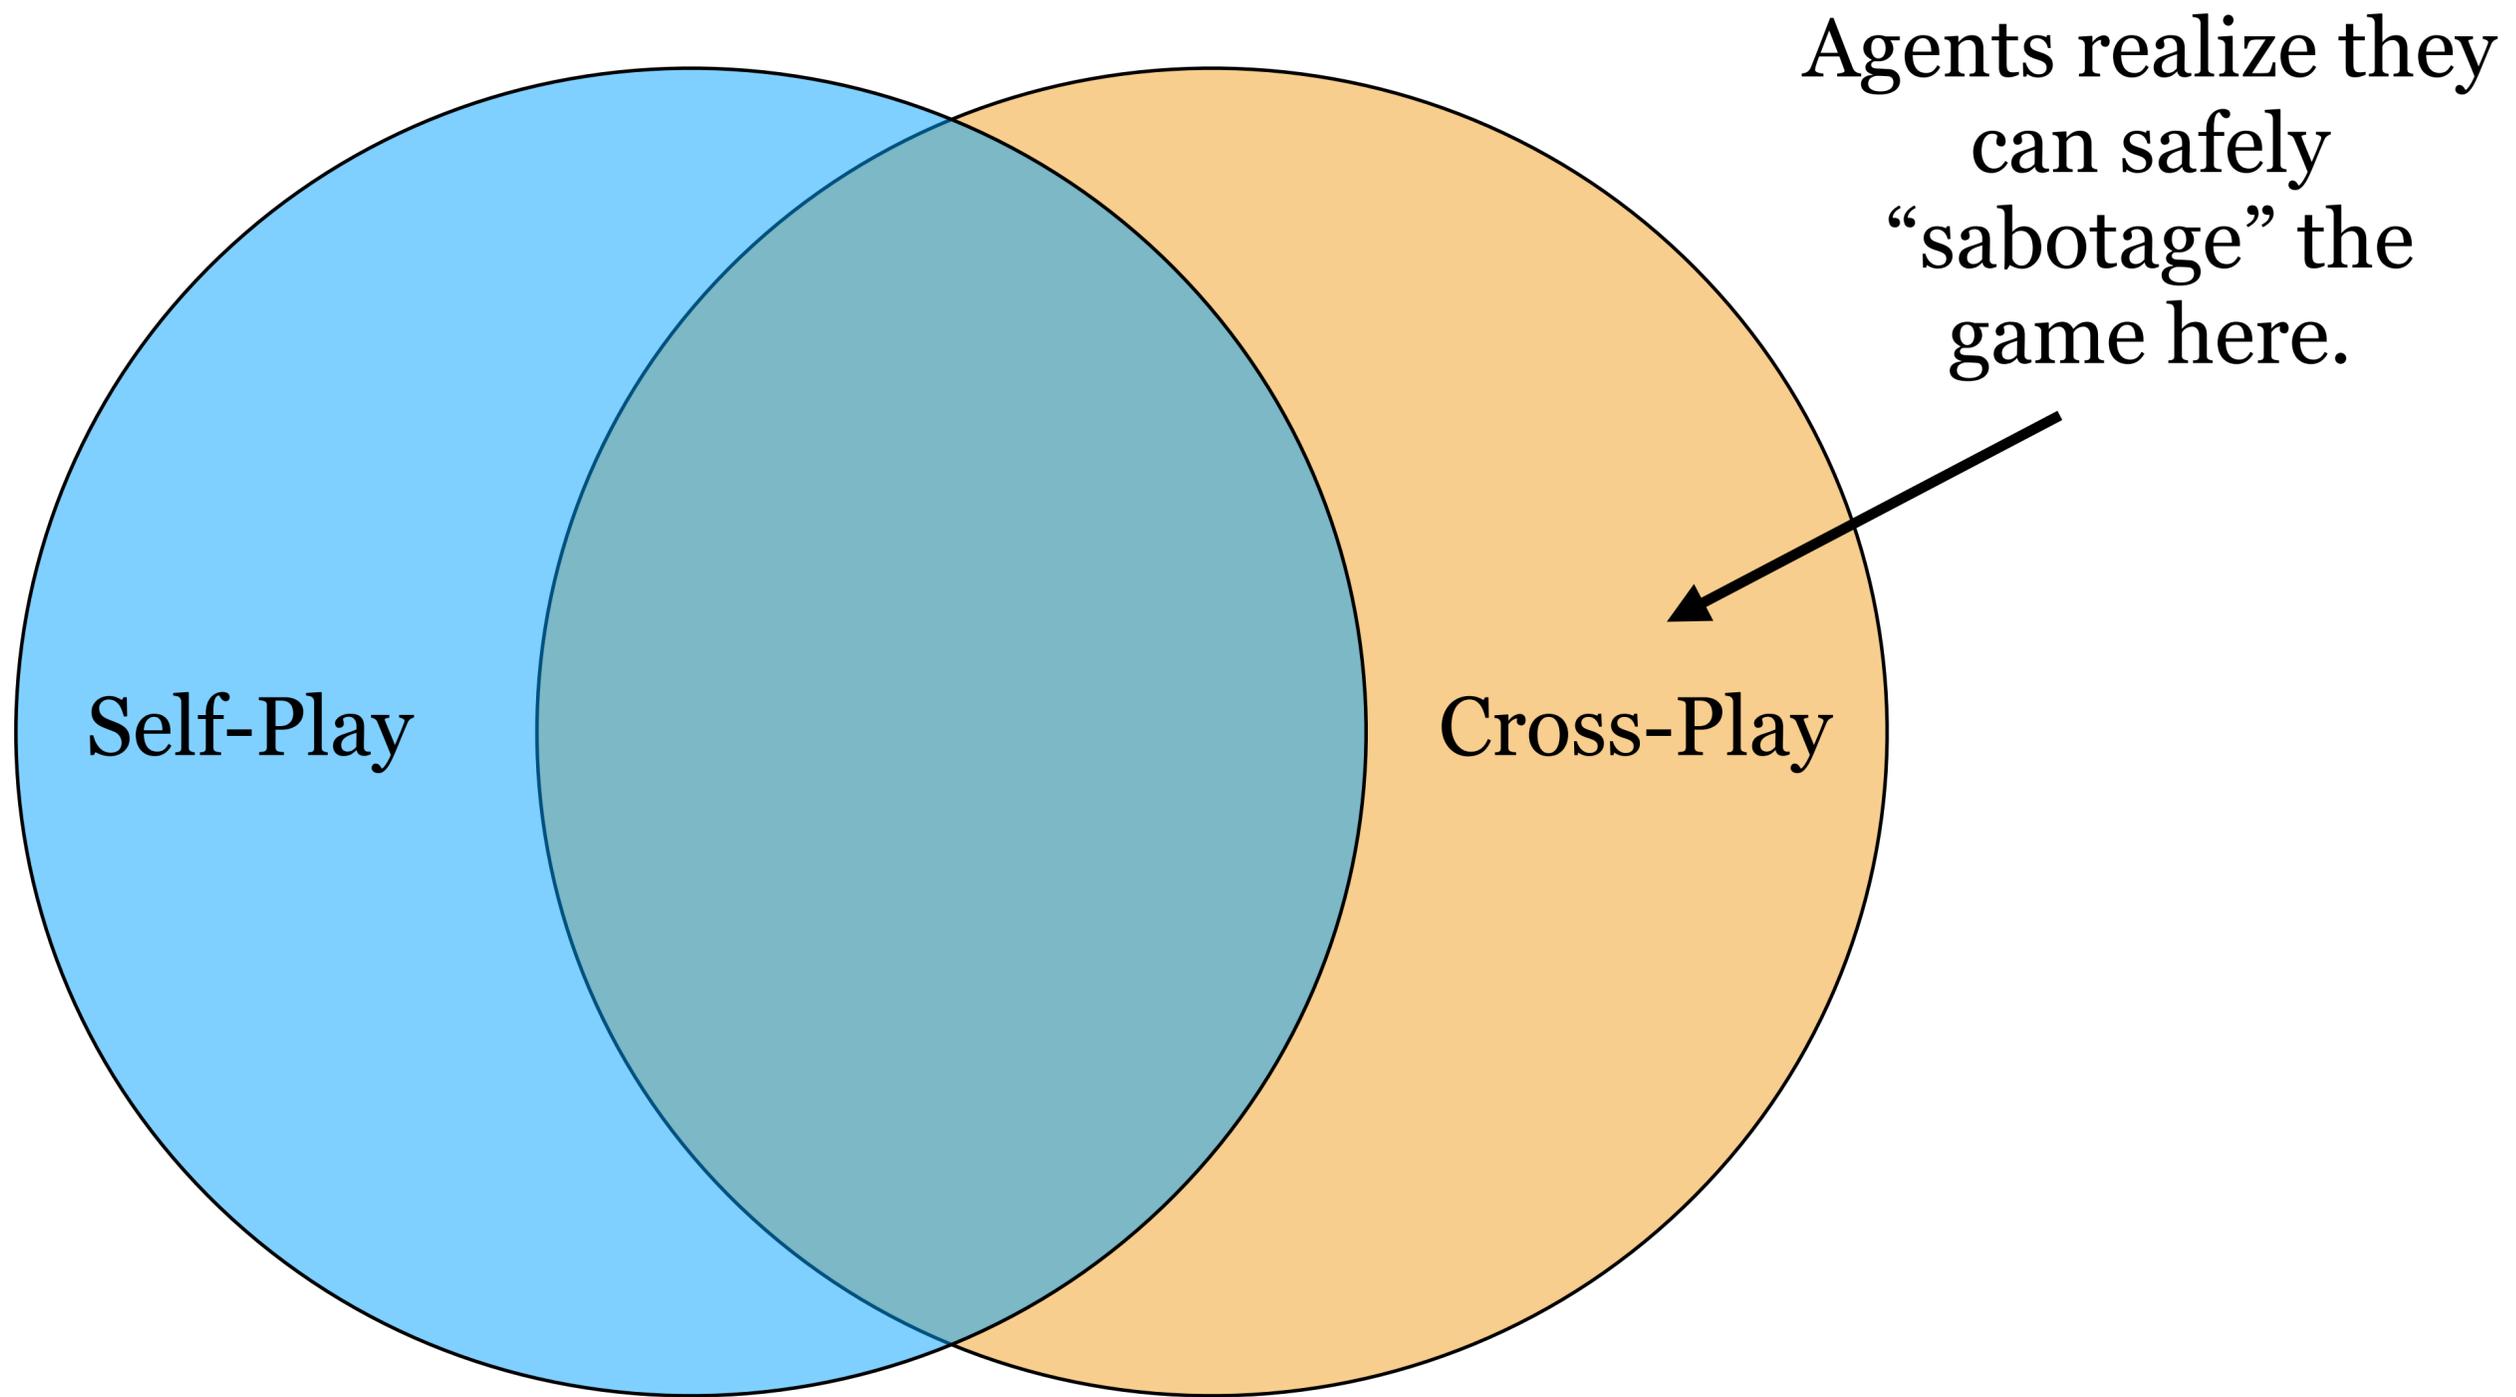


Cross-play  
incorrectly values  
this convention!



**Self-Play**

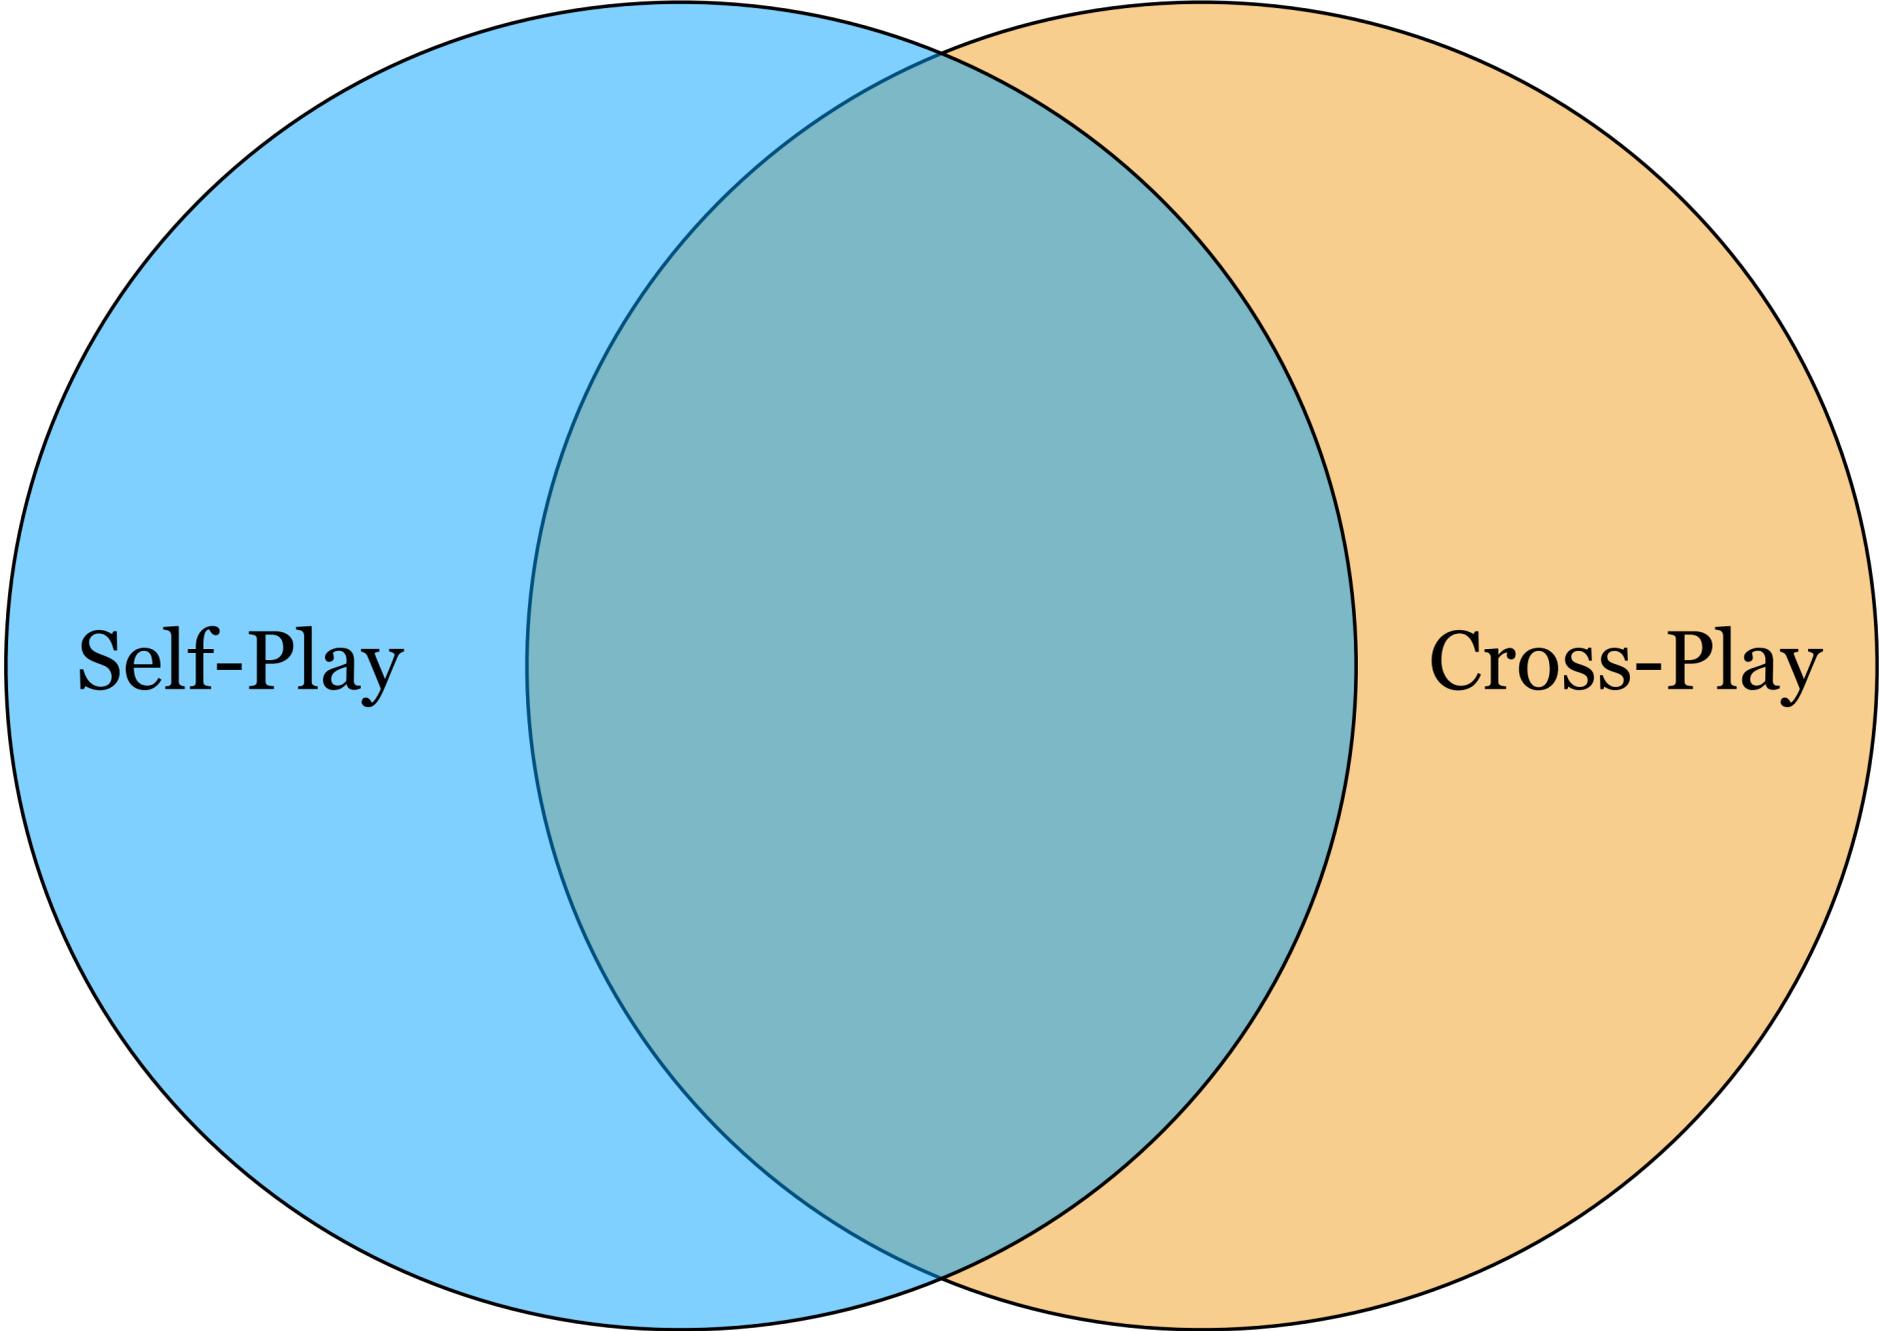
**Cross-Play**



Self-Play

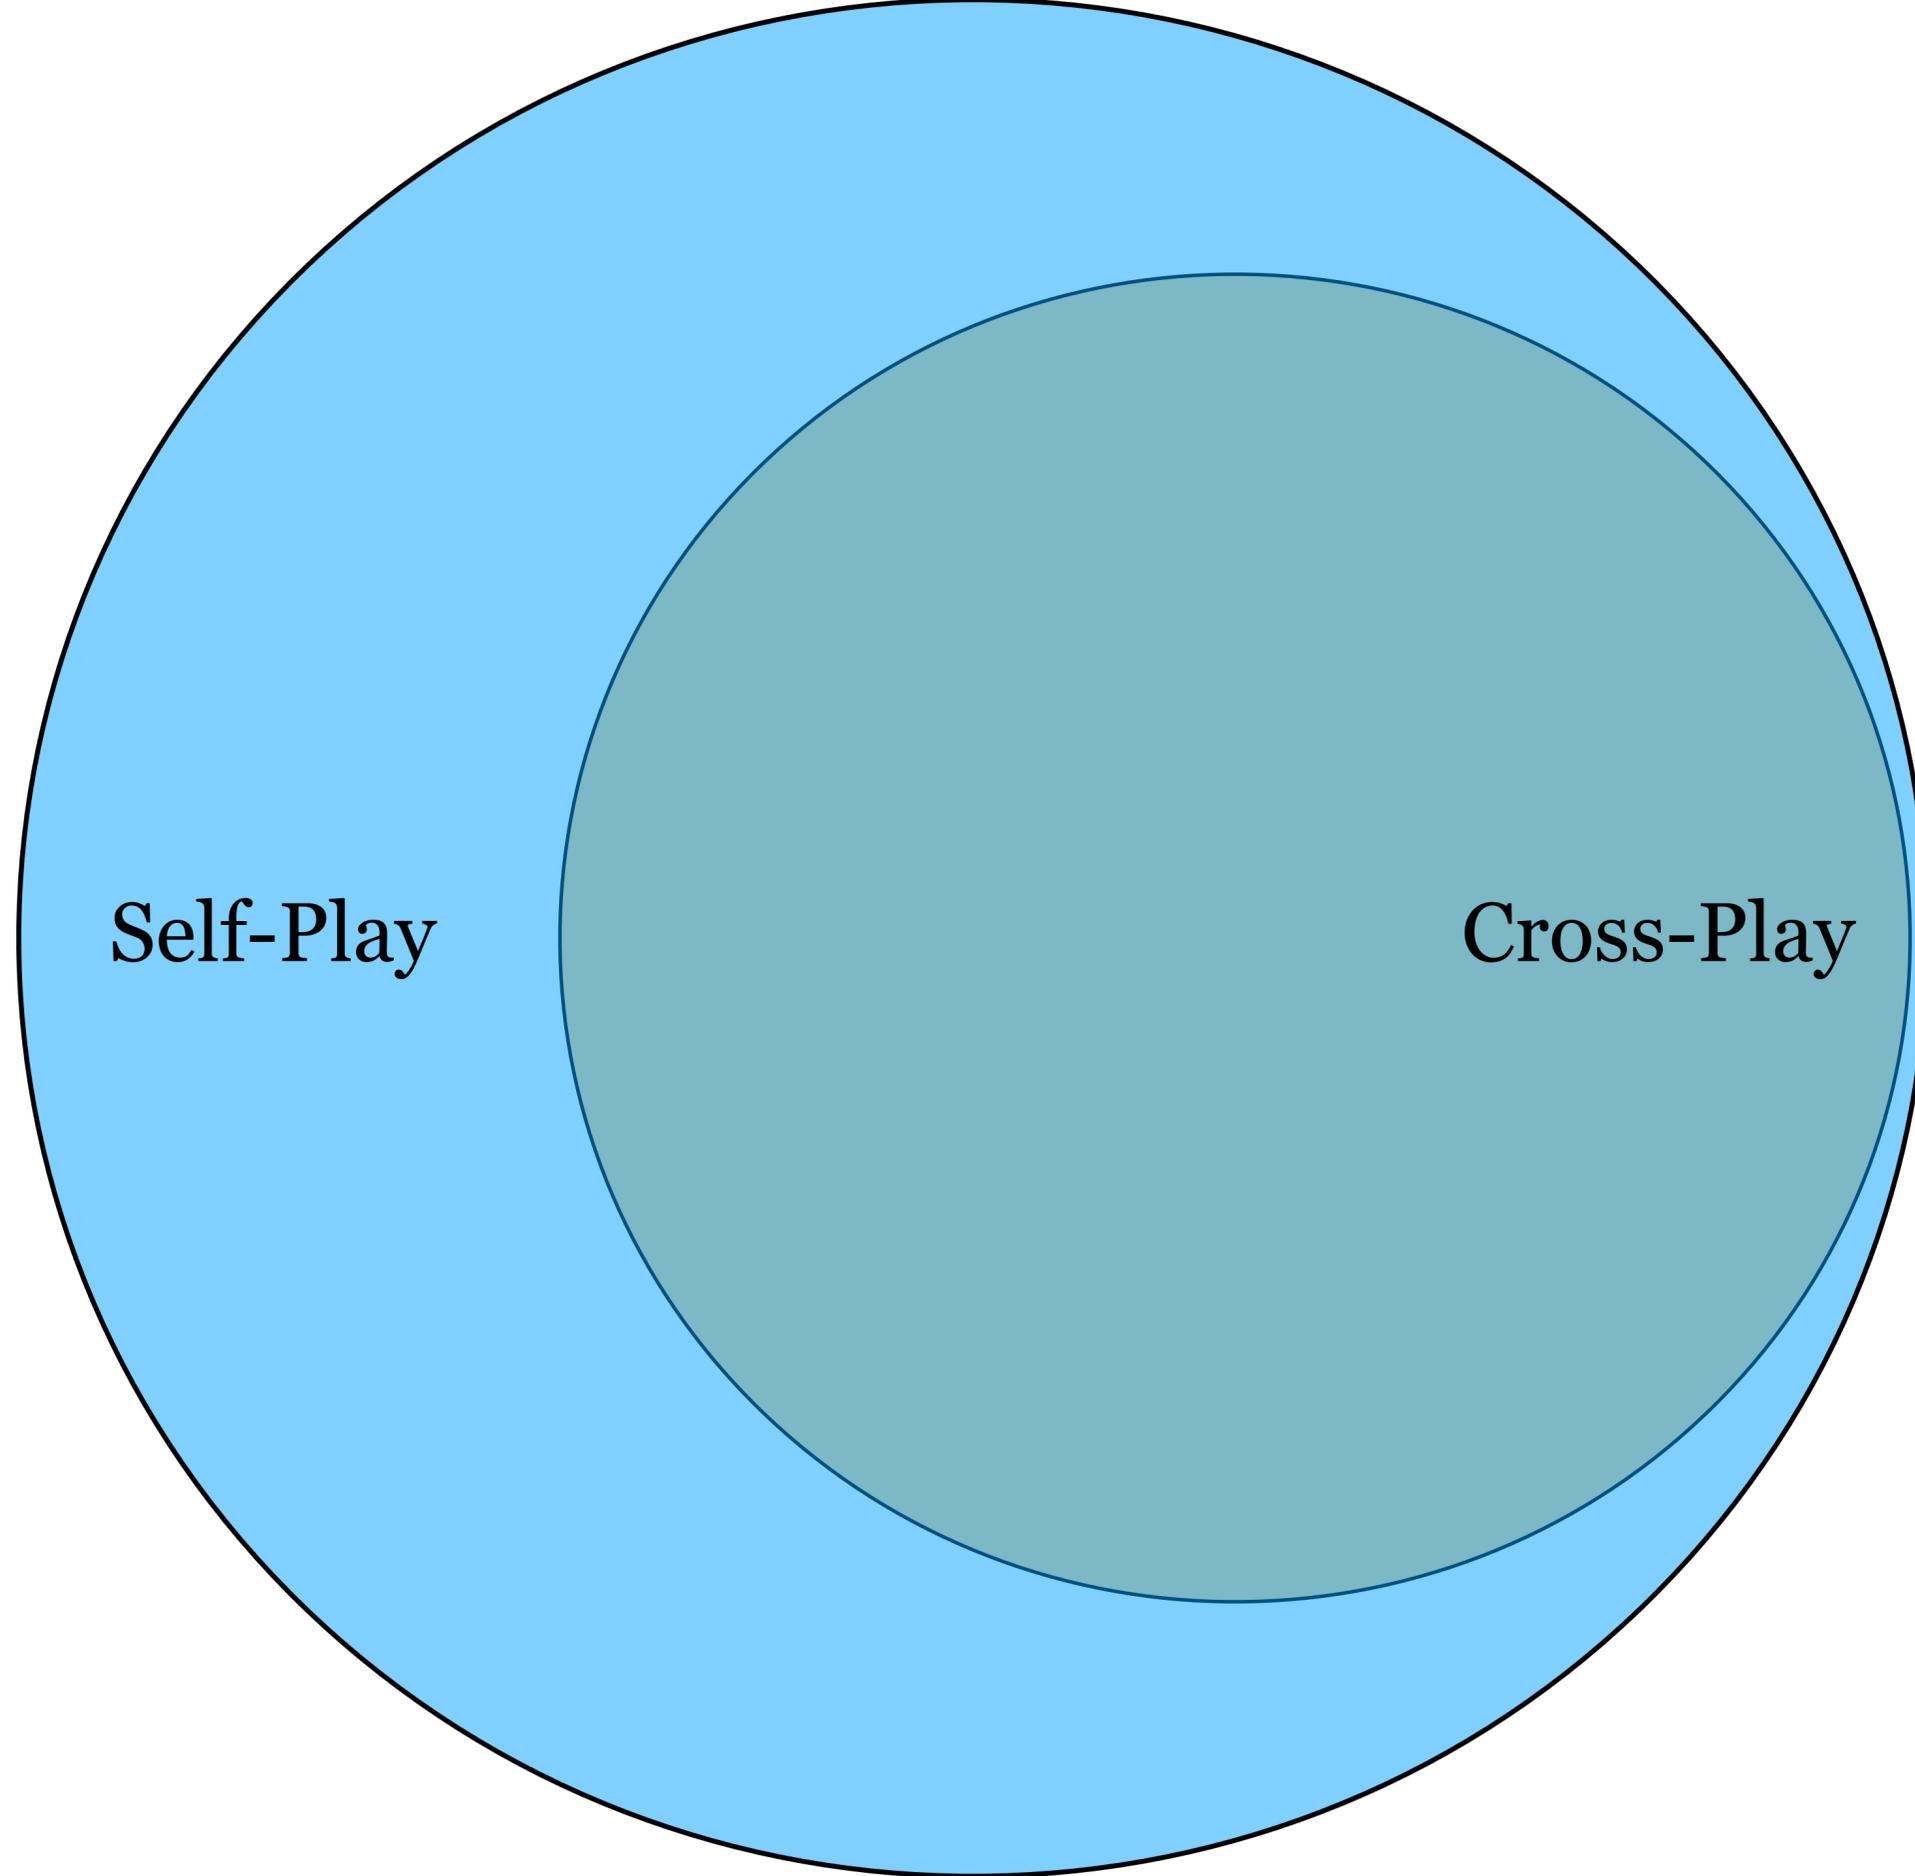
Cross-Play

Agents realize they  
can safely  
"sabotage" the  
game here.



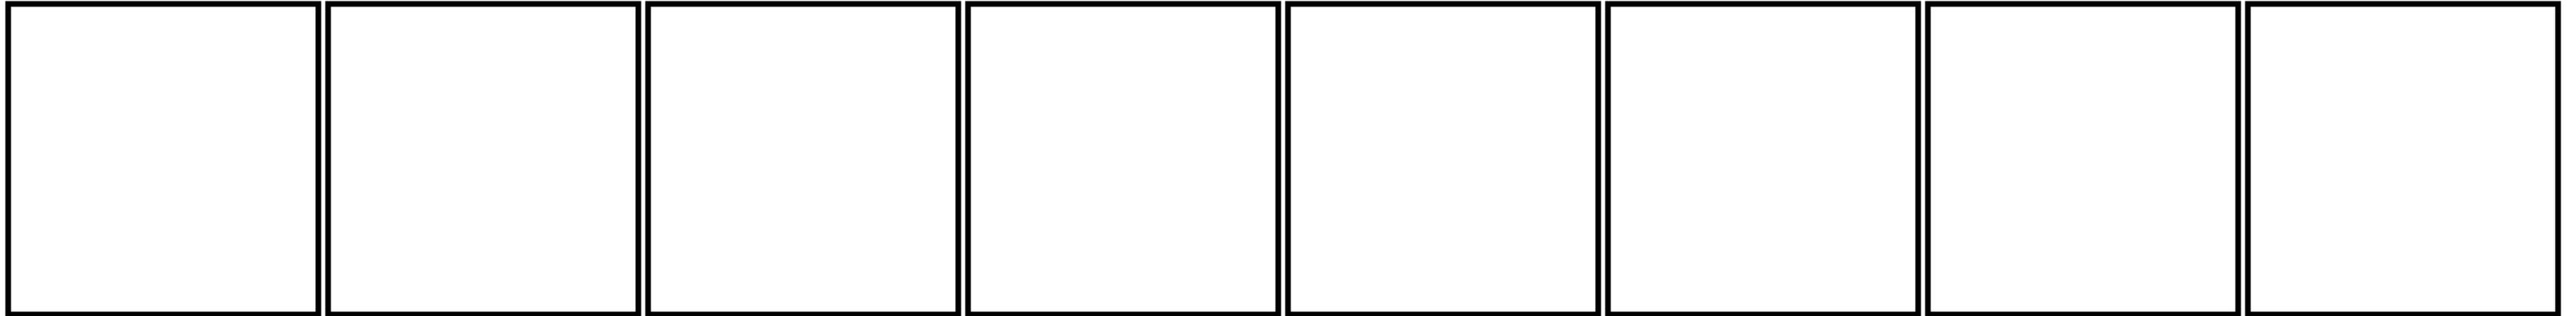
**Self-Play**

**Cross-Play**



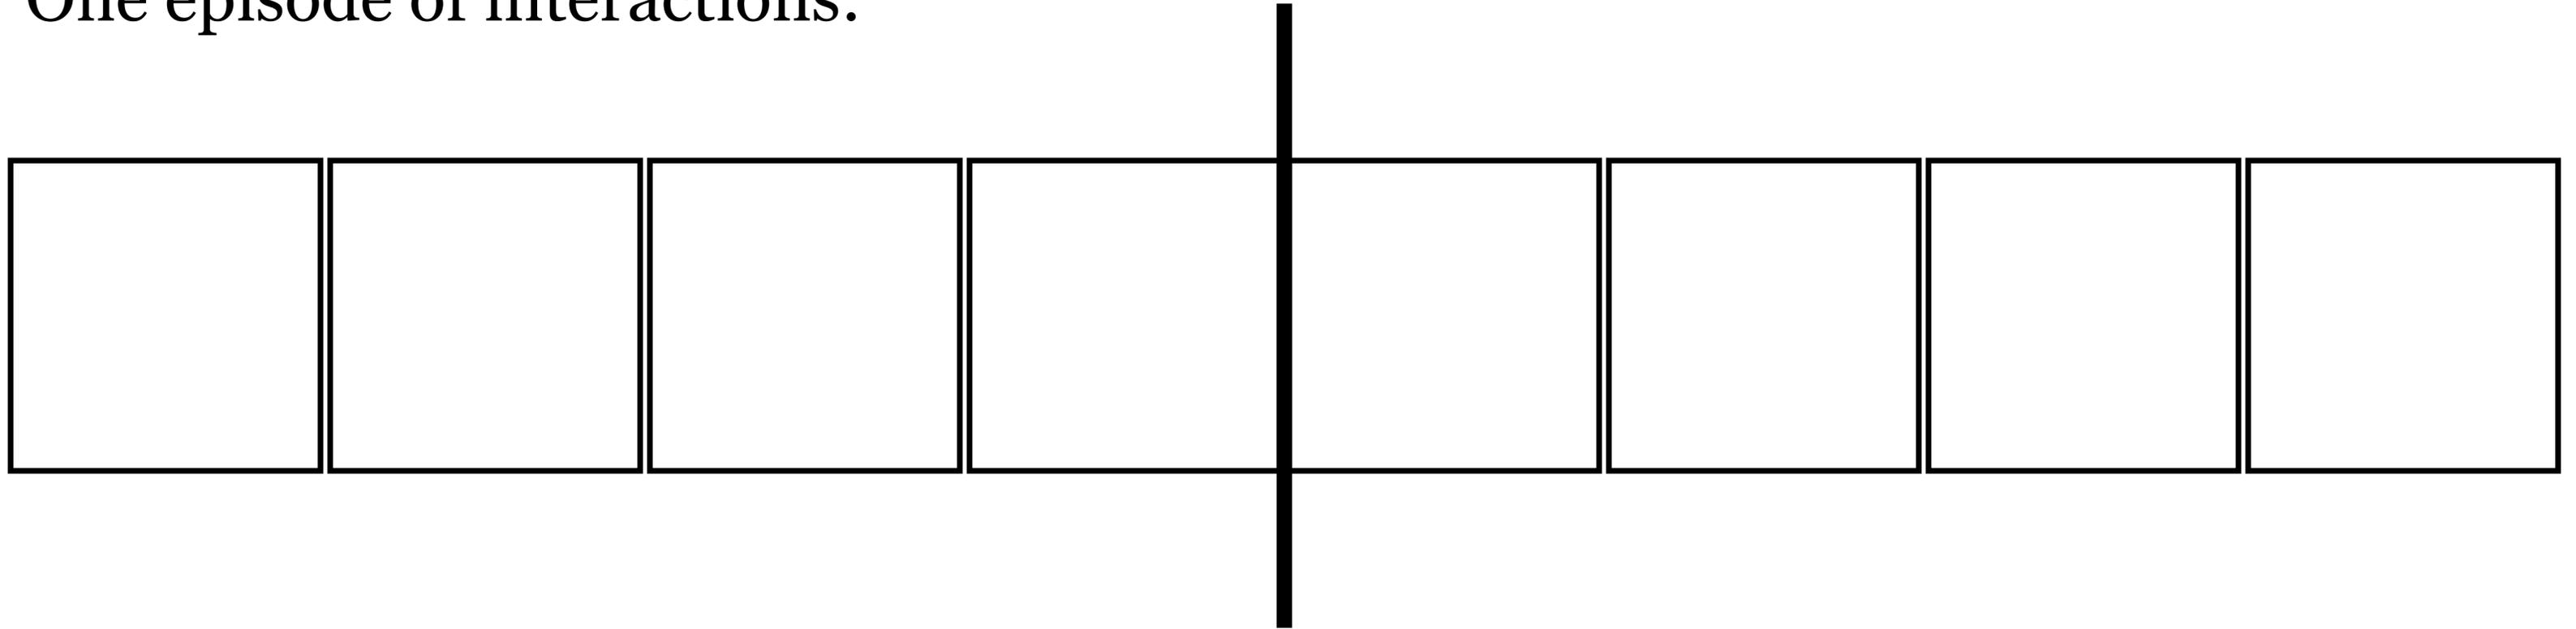
# Mixed-Play

One episode of interactions:



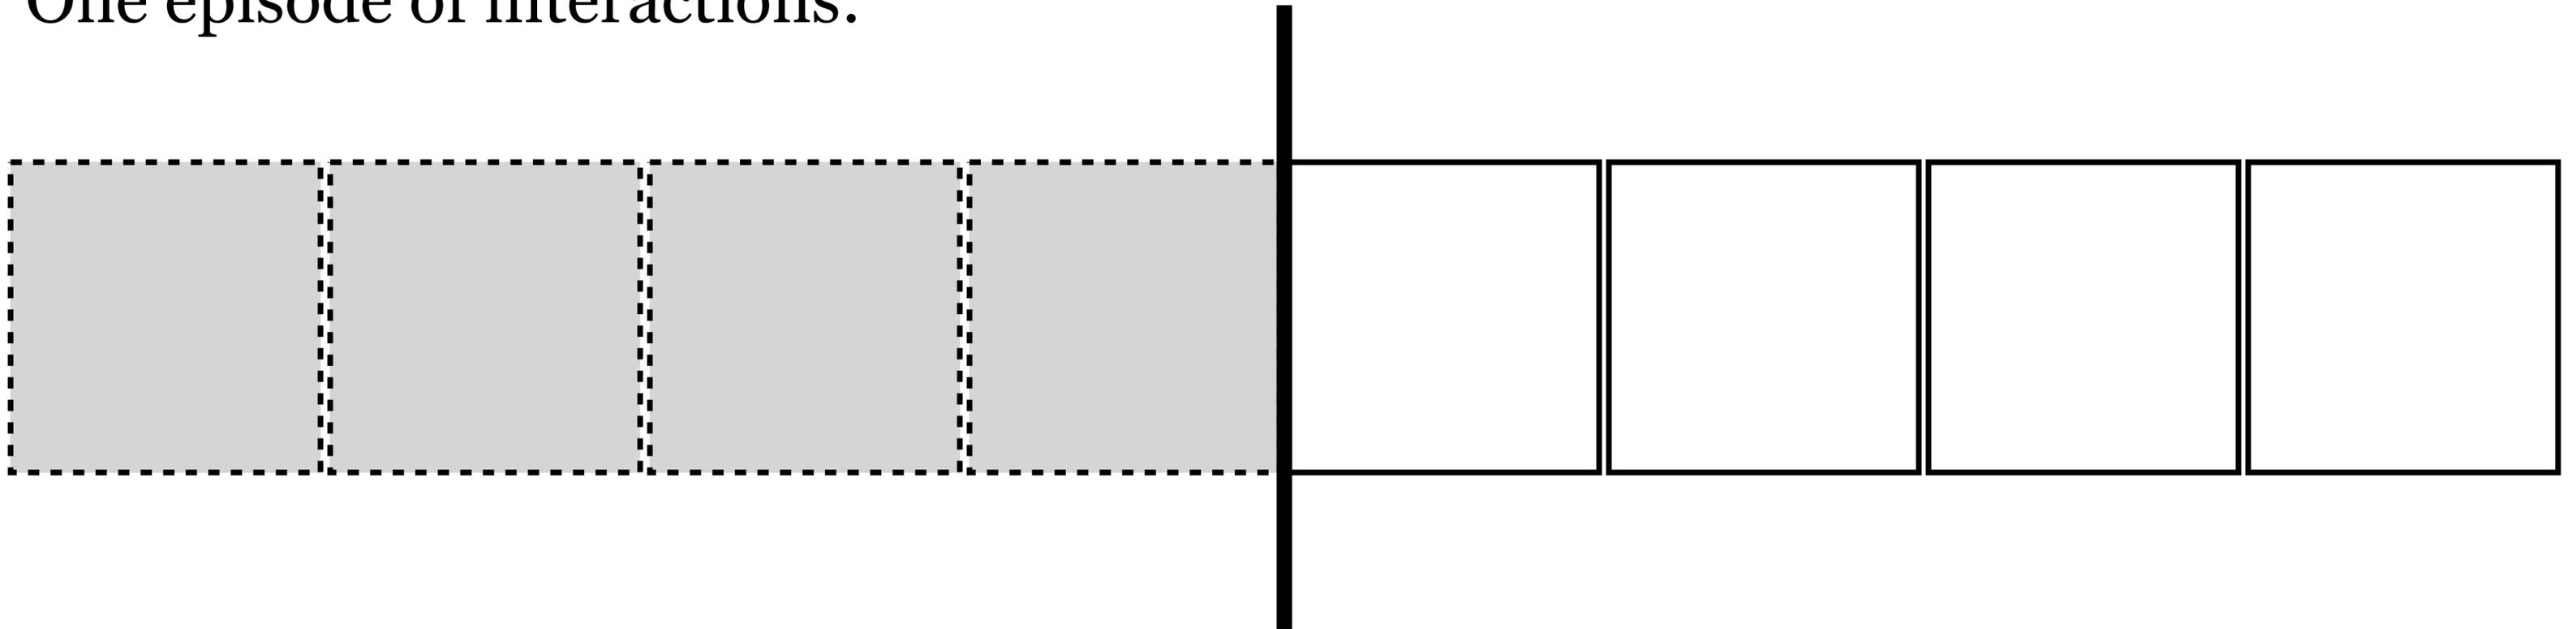
# Mixed-Play

One episode of interactions:



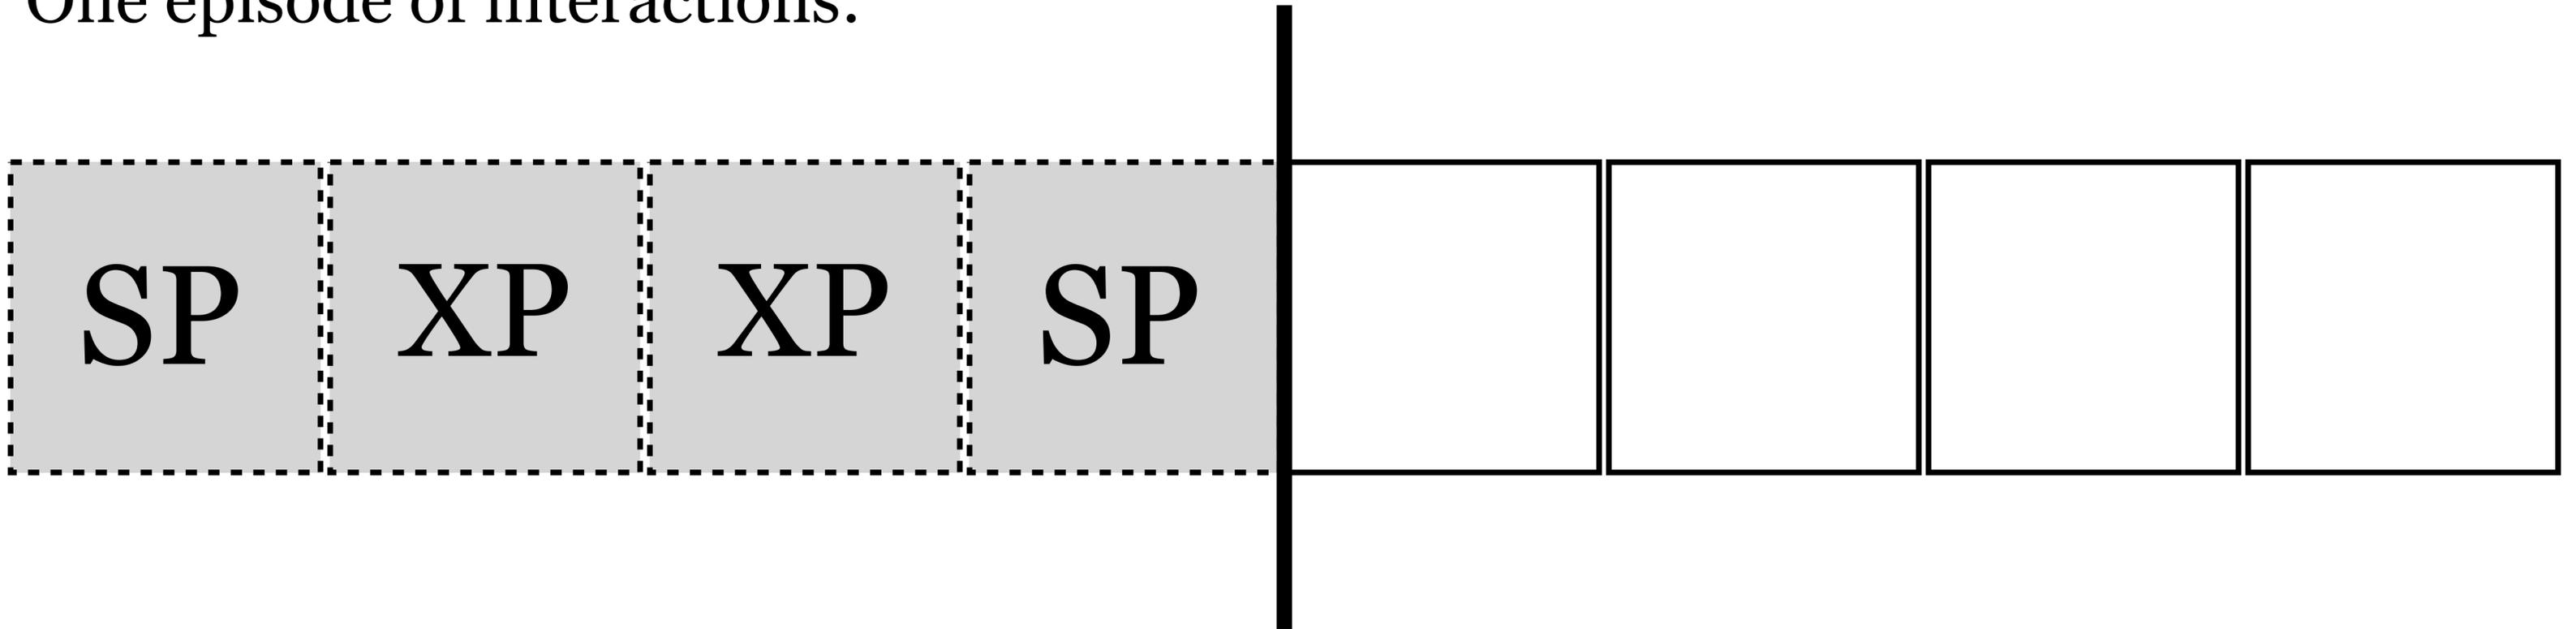
# Mixed-Play

One episode of interactions:



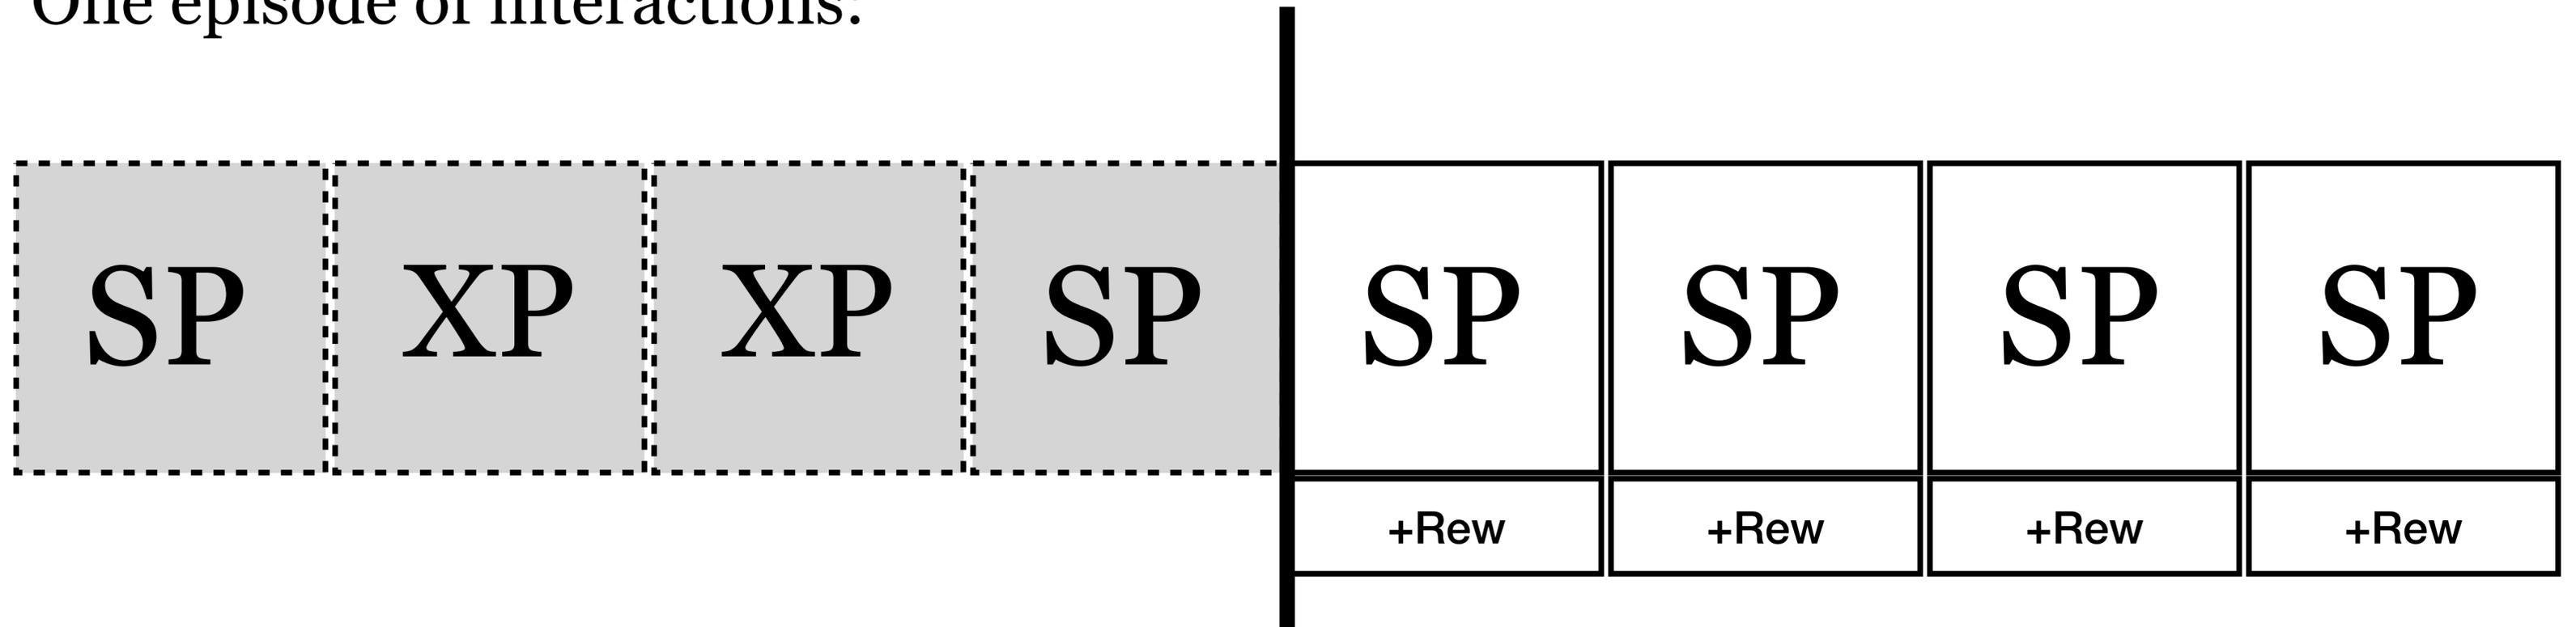
# Mixed-Play

One episode of interactions:

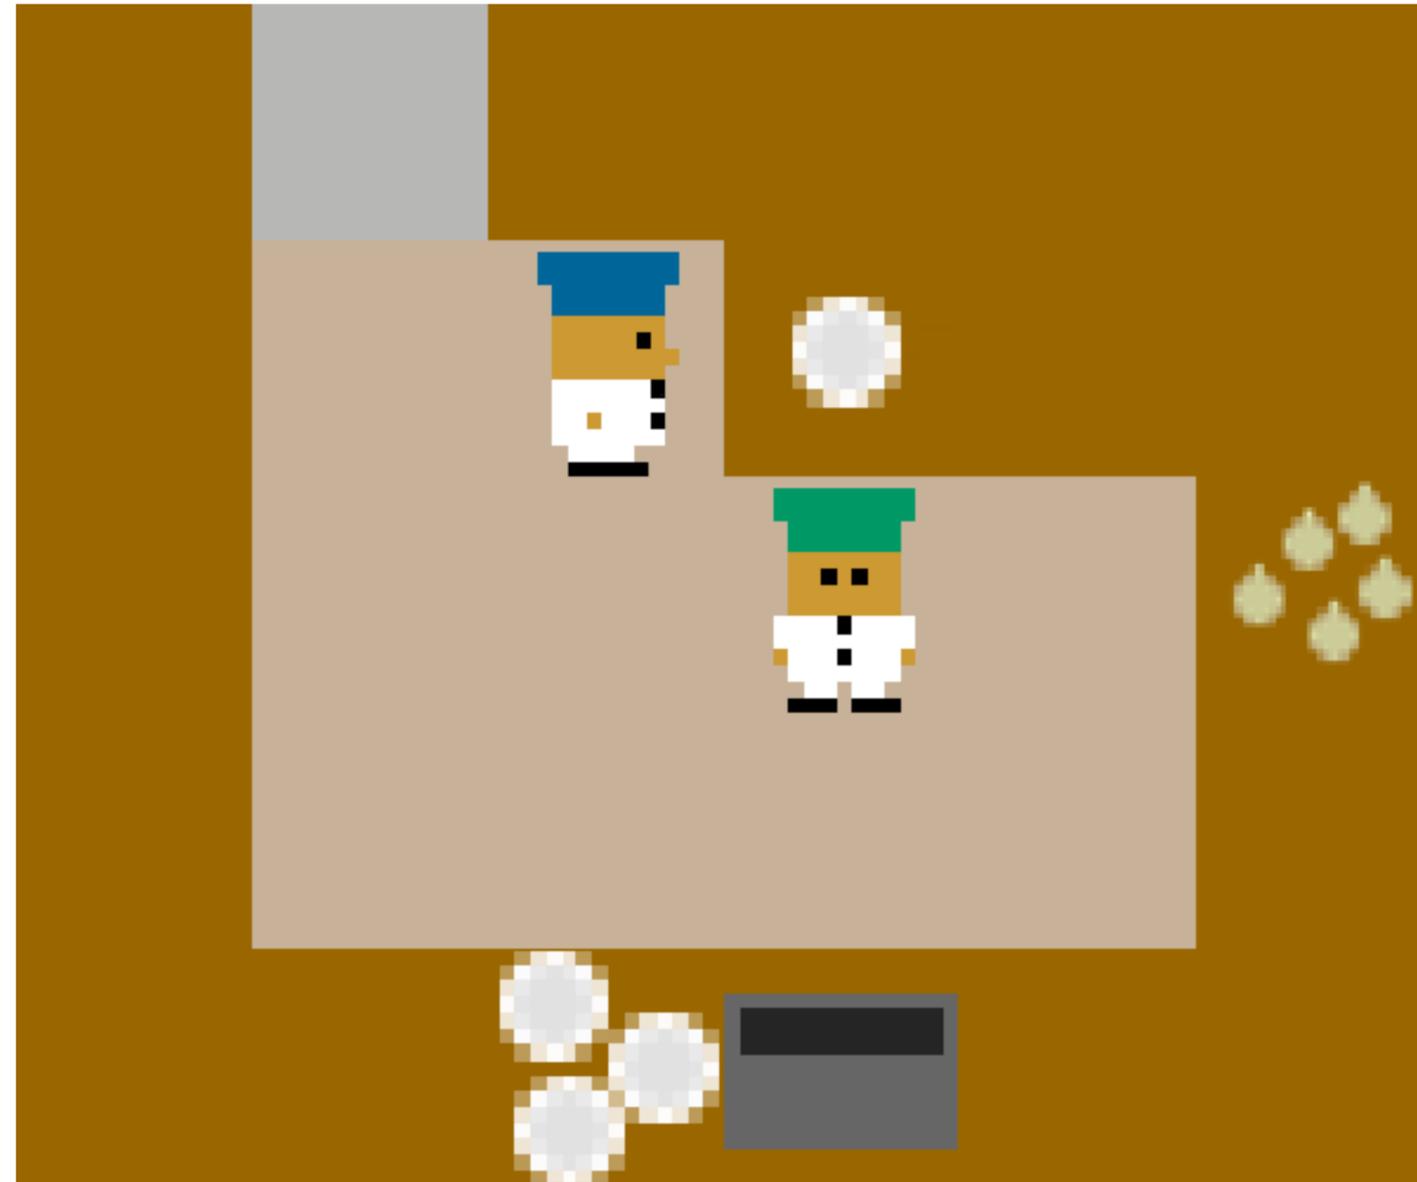


# Mixed-Play

One episode of interactions:



# Mixed-Play



# Mixed-Play



Mixed-Play fixes  
handshakes!

# CoMeDi

Minimize loss function to generate the  $n^{\text{th}}$  convention:

$$L(\pi_n) = -J(\pi_n, \pi_n) + \alpha J(\pi_n, \pi^*) - \beta J_M(\pi_n, \pi^*)$$

$\pi^*$  is the most-compatible previously discovered convention

$J_M$  is the expected return from the self-play phase of mixed-play

# CoMeDi

Minimize loss function to generate the  $n^{\text{th}}$  convention:

$$L(\pi_n) = -J(\pi_n, \pi_n) + \alpha J(\pi_n, \pi^*) - \beta J_M(\pi_n, \pi^*)$$

↑  
Maximize  
Self-Play

$\pi^*$  is the most-compatible previously discovered convention

$J_M$  is the expected return from the self-play phase of mixed-play

# CoMeDi

Minimize loss function to generate the  $n^{\text{th}}$  convention:

$$L(\pi_n) = -J(\pi_n, \pi_n) + \alpha J(\pi_n, \pi^*) - \beta J_M(\pi_n, \pi^*)$$

↑  
Maximize  
Self-Play

↑  
Minimize  
Cross-Play

$\pi^*$  is the most-compatible previously discovered convention

$J_M$  is the expected return from the self-play phase of mixed-play

# CoMeDi

Minimize loss function to generate the  $n^{\text{th}}$  convention:

$$L(\pi_n) = -J(\pi_n, \pi_n) + \alpha J(\pi_n, \pi^*) - \beta J_M(\pi_n, \pi^*)$$

↑  
Maximize  
Self-Play

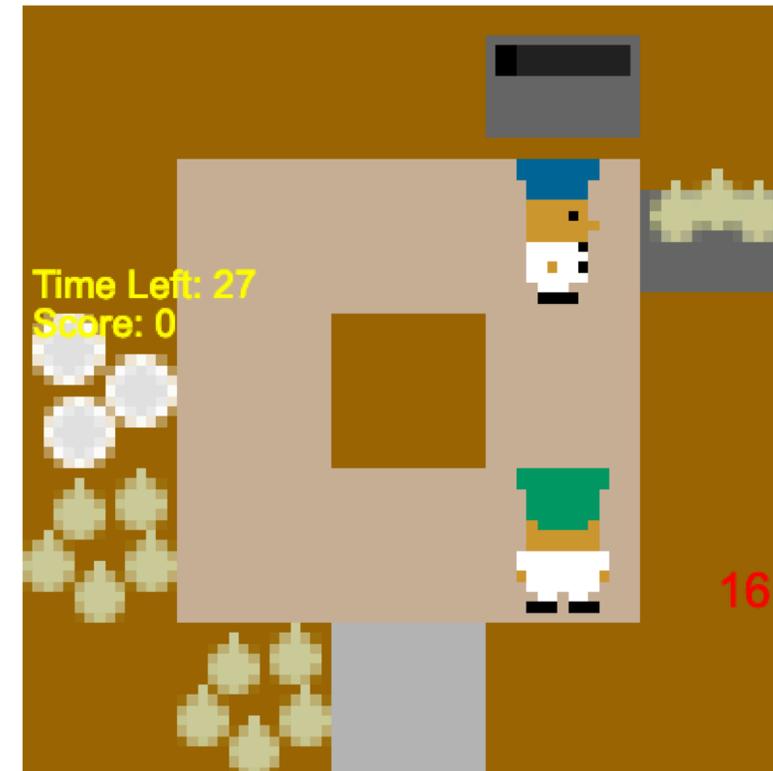
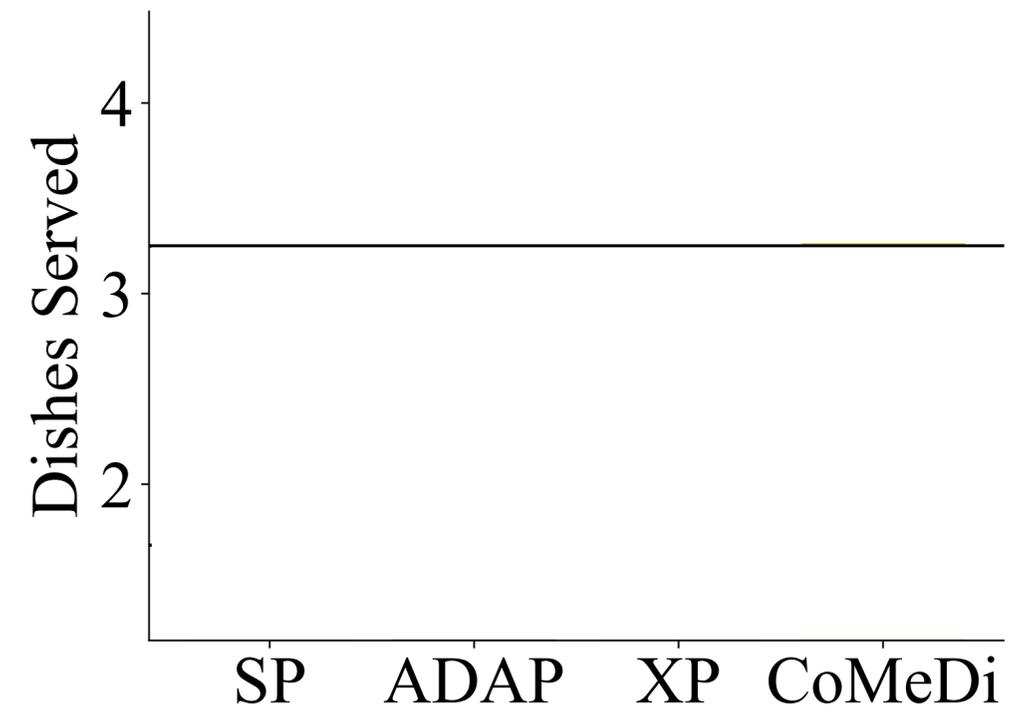
↑  
Minimize  
Cross-Play

↑  
Maximize  
Mixed-Play

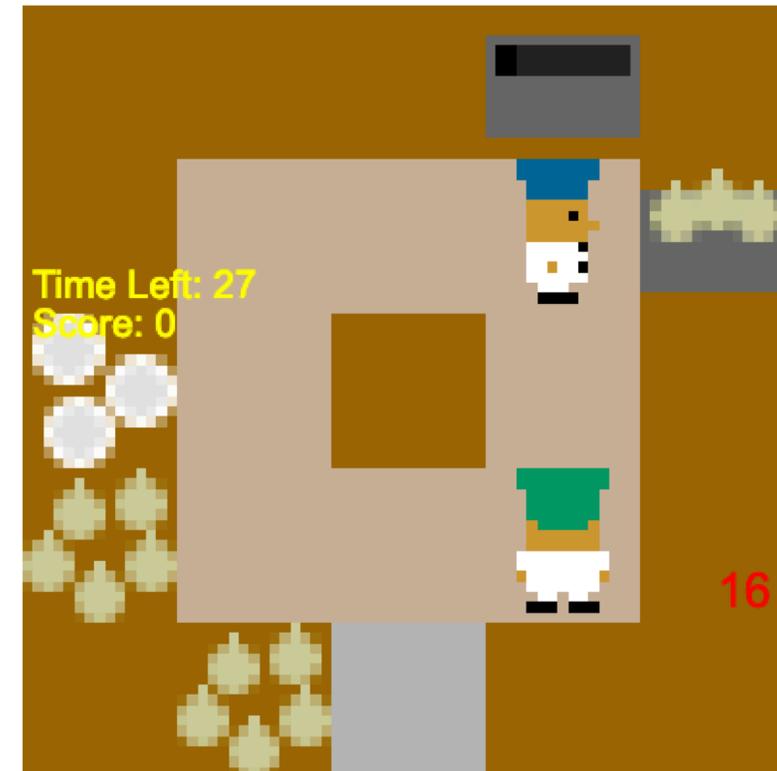
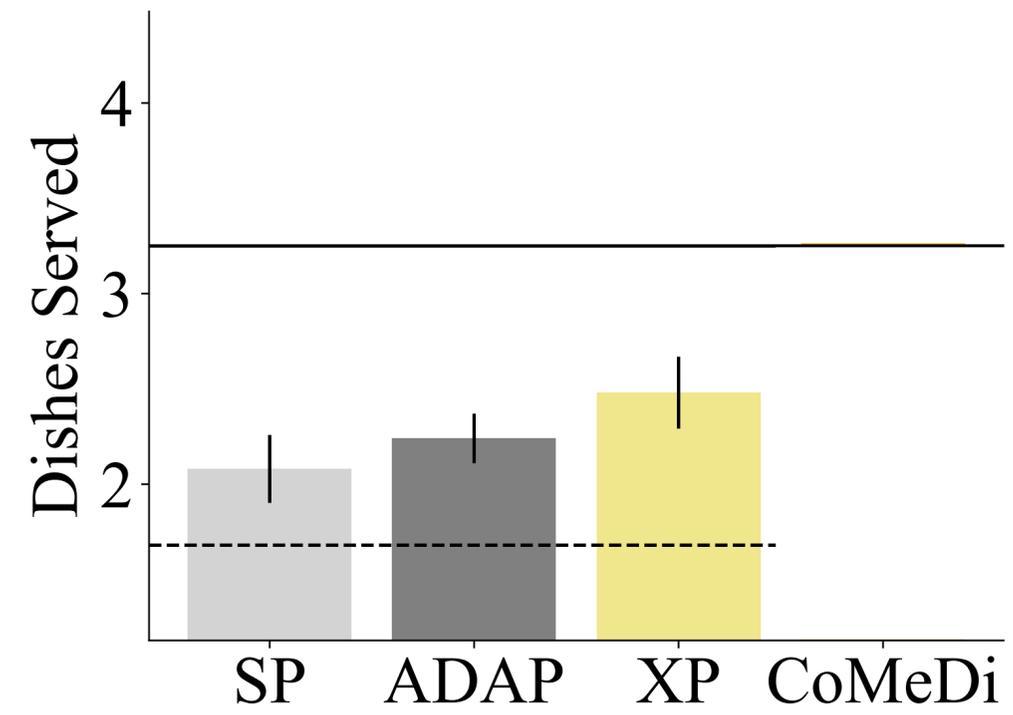
$\pi^*$  is the most-compatible previously discovered convention

$J_M$  is the expected return from the self-play phase of mixed-play

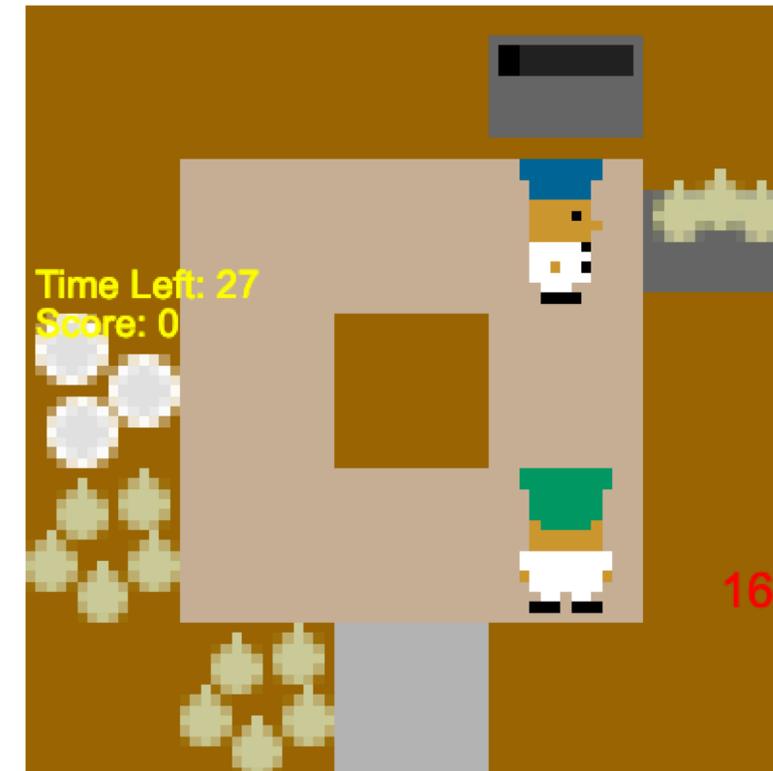
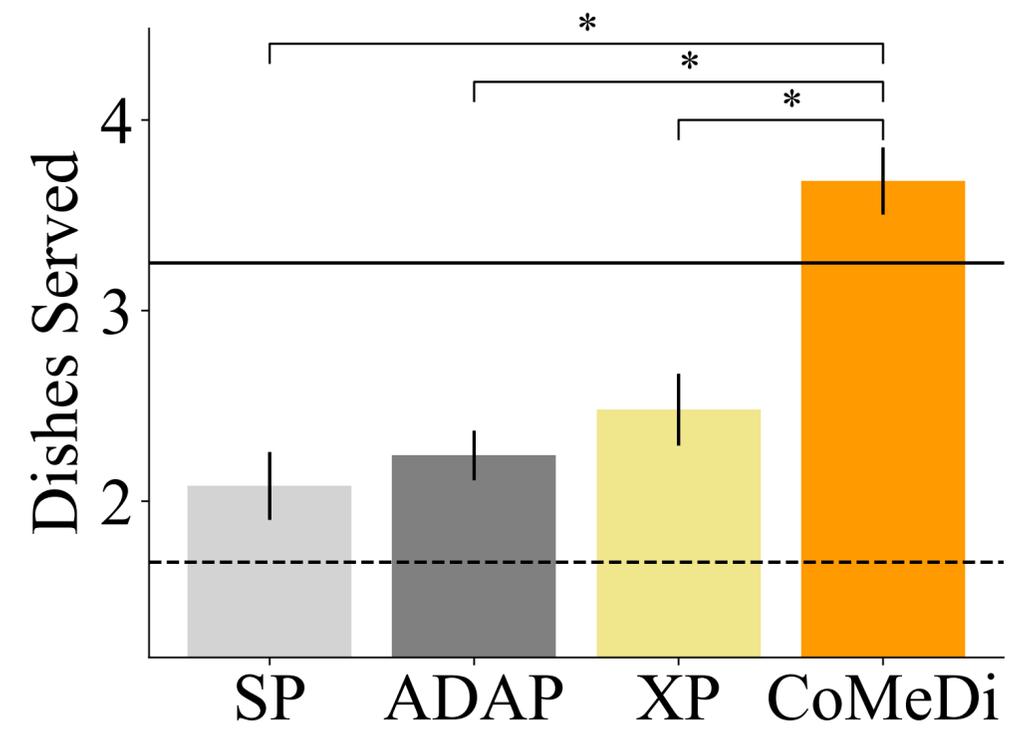
# User Study



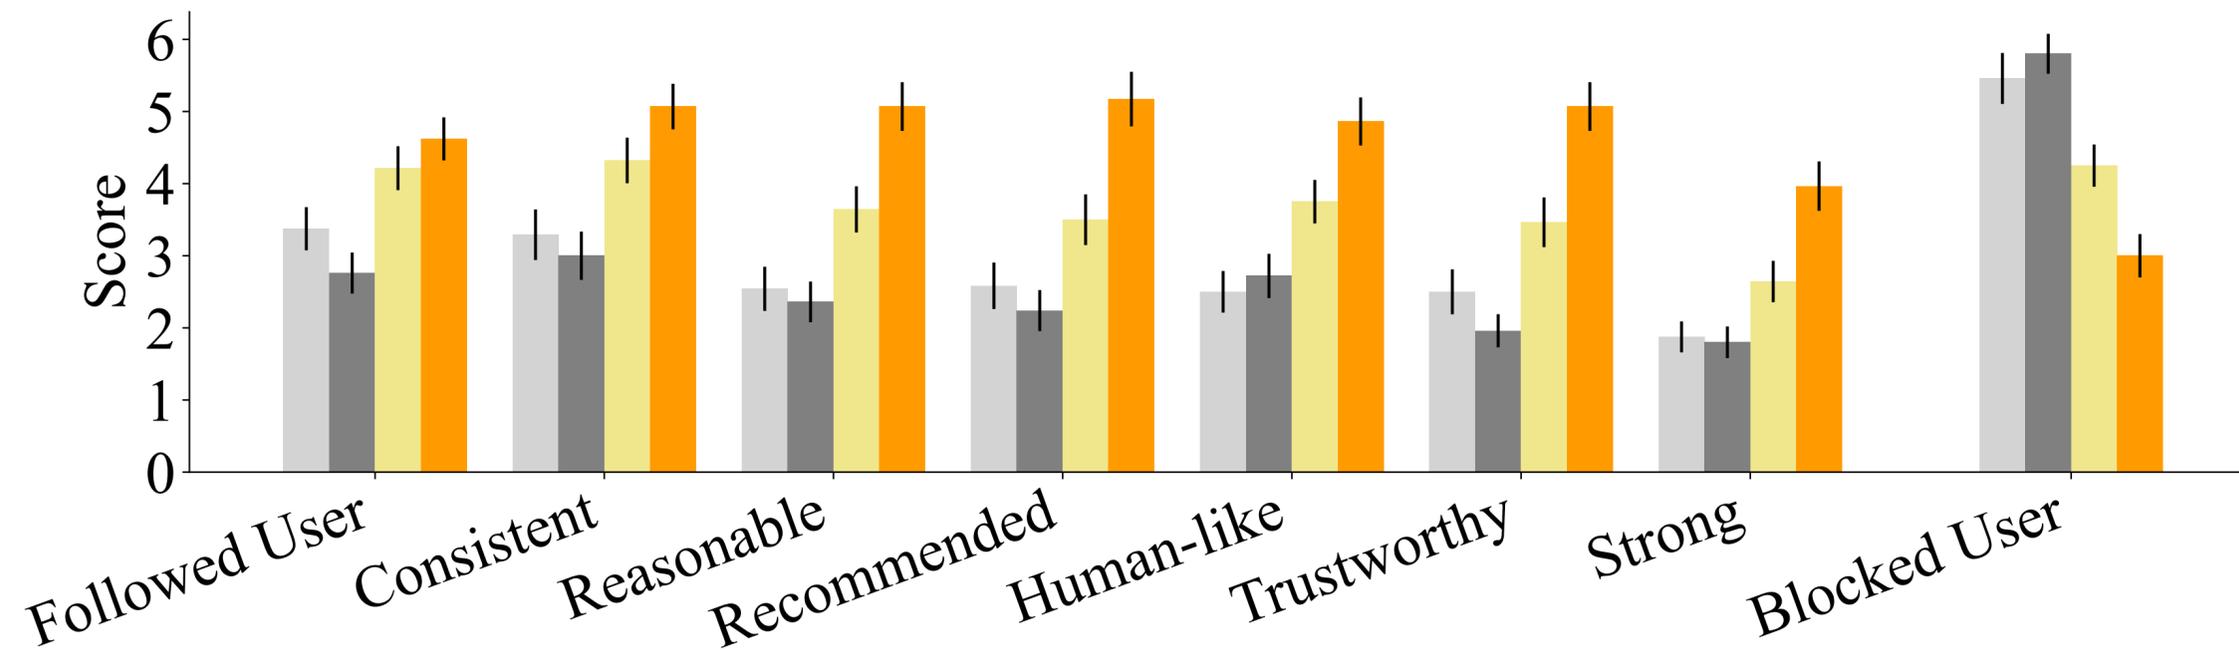
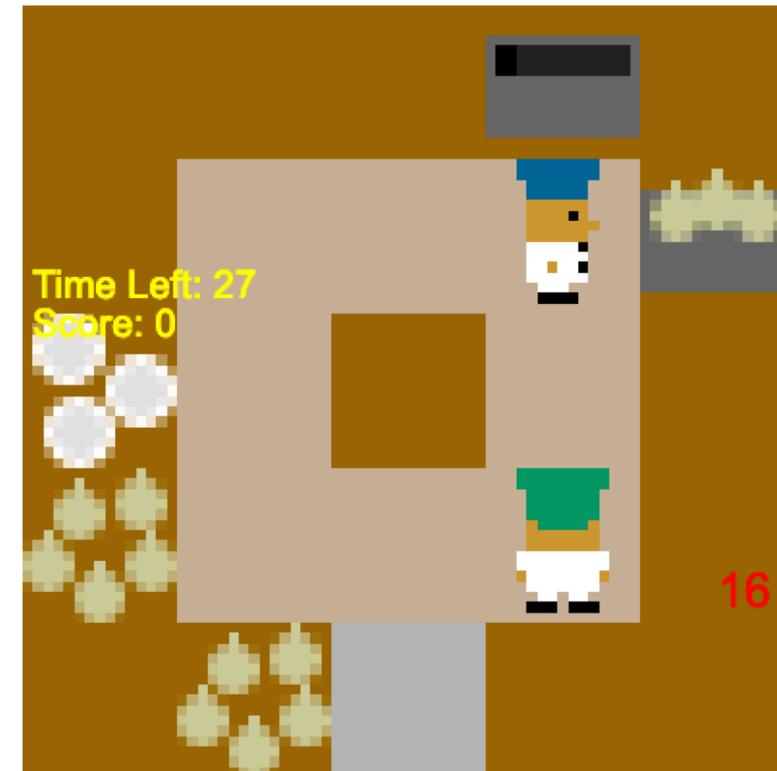
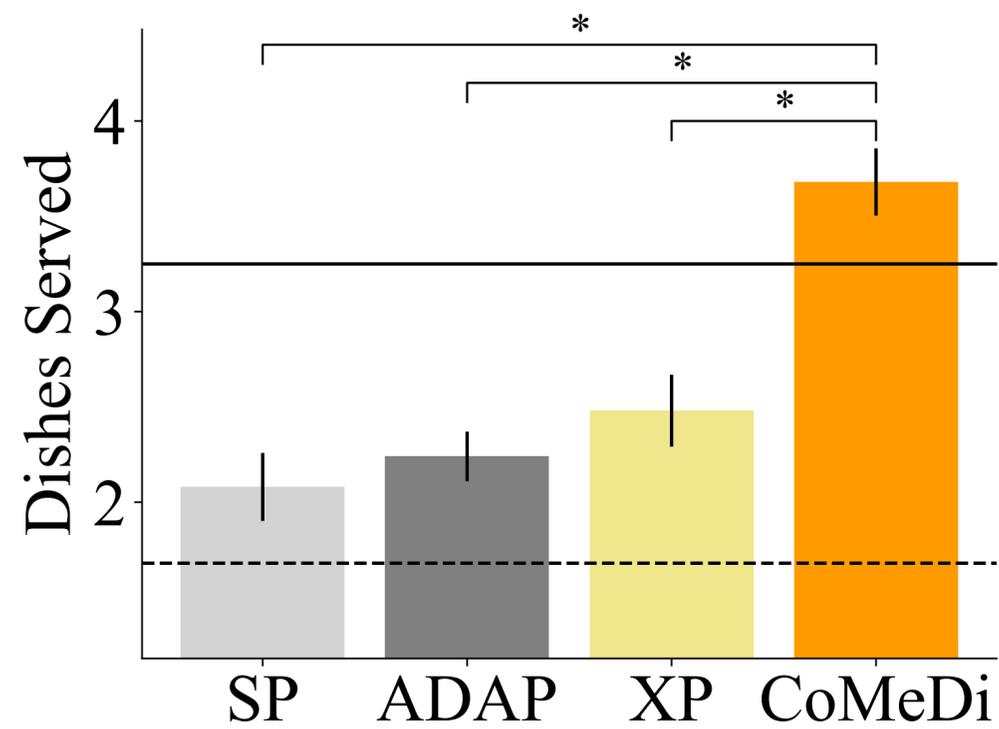
# User Study



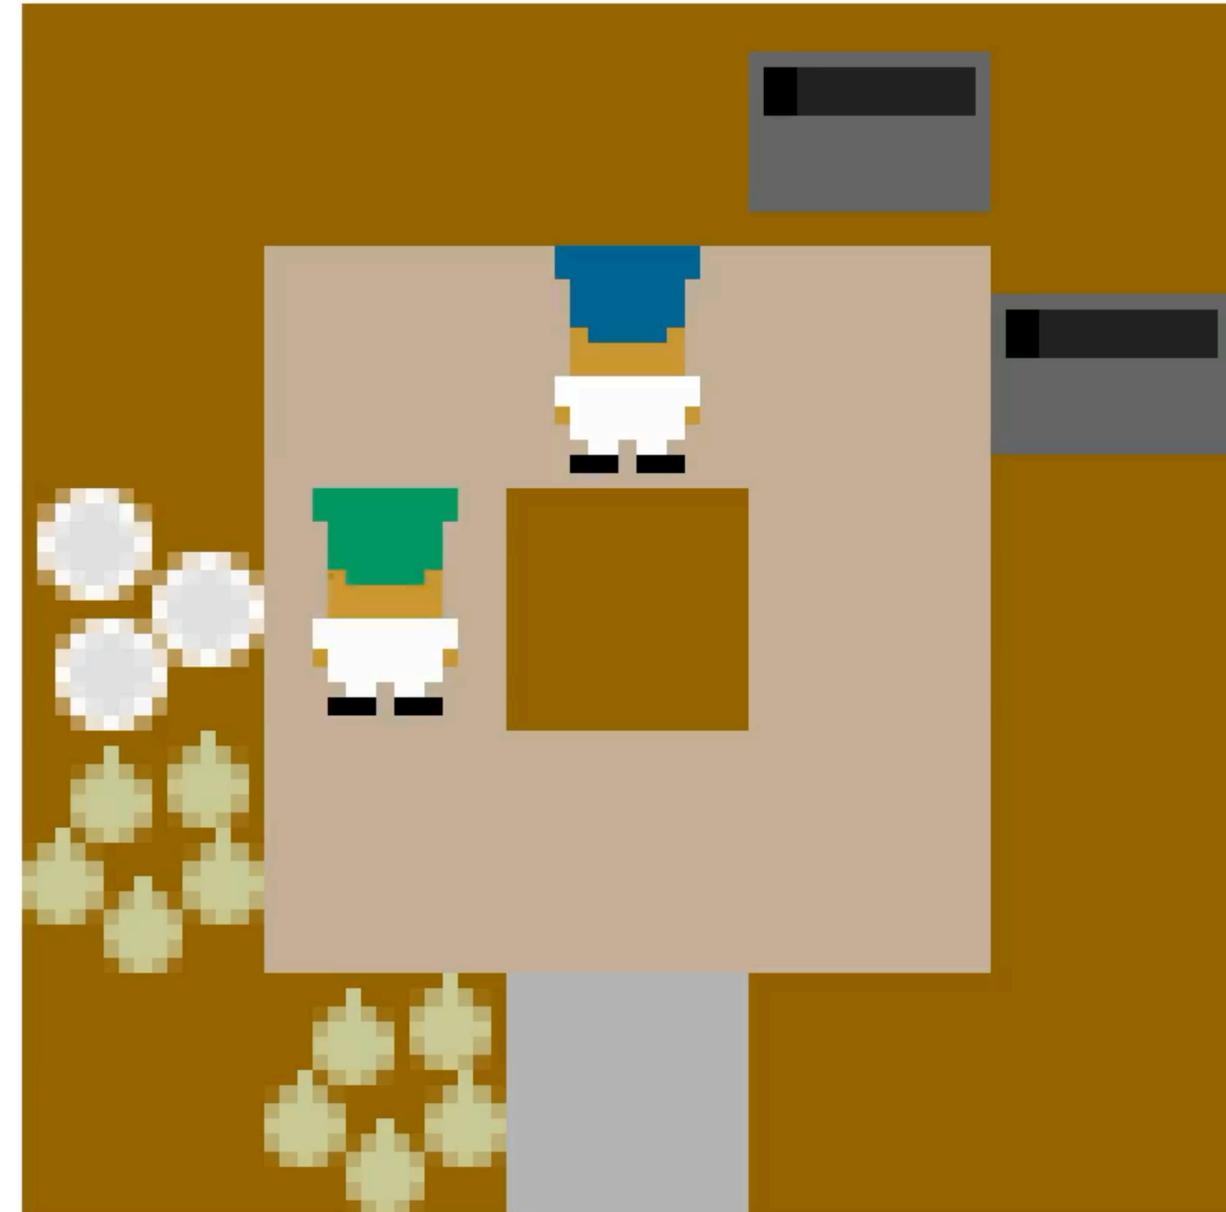
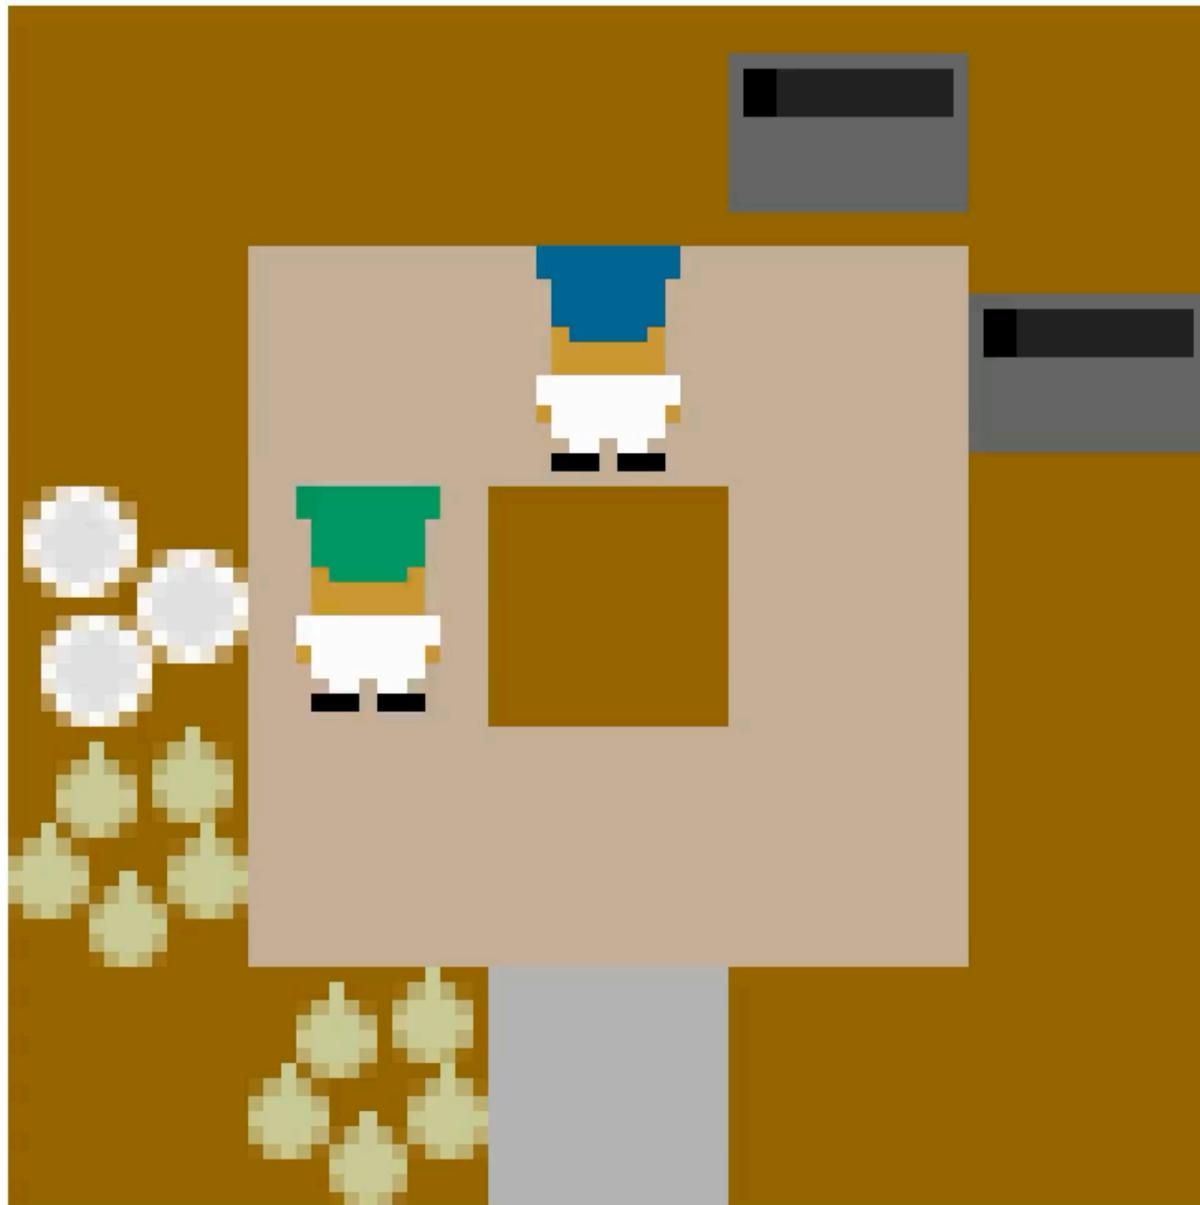
# User Study



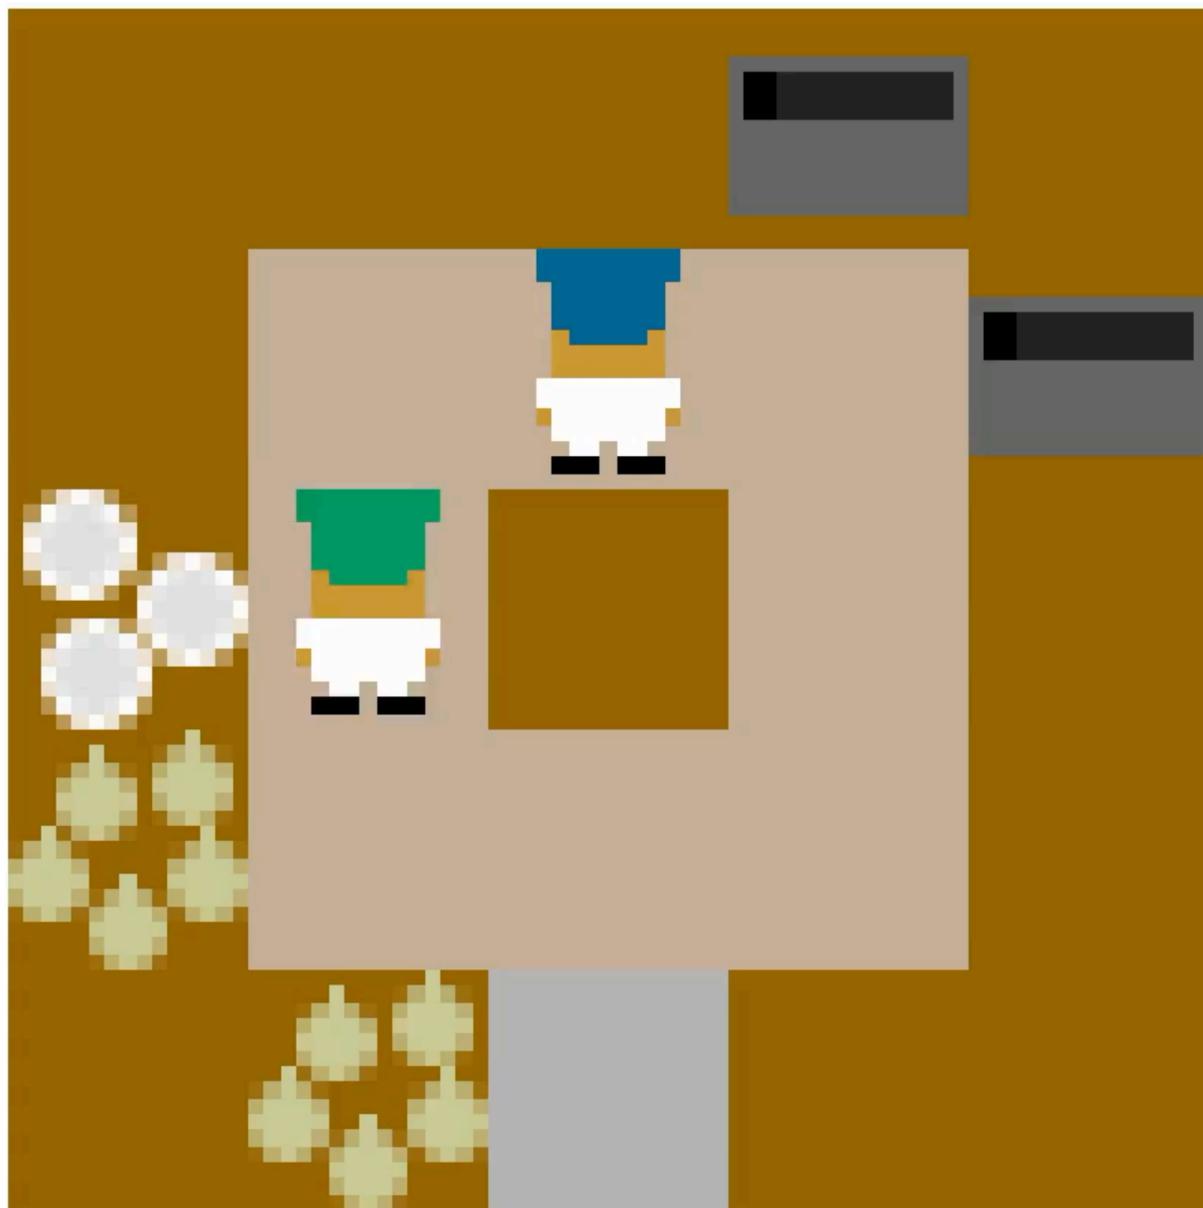
# User Study



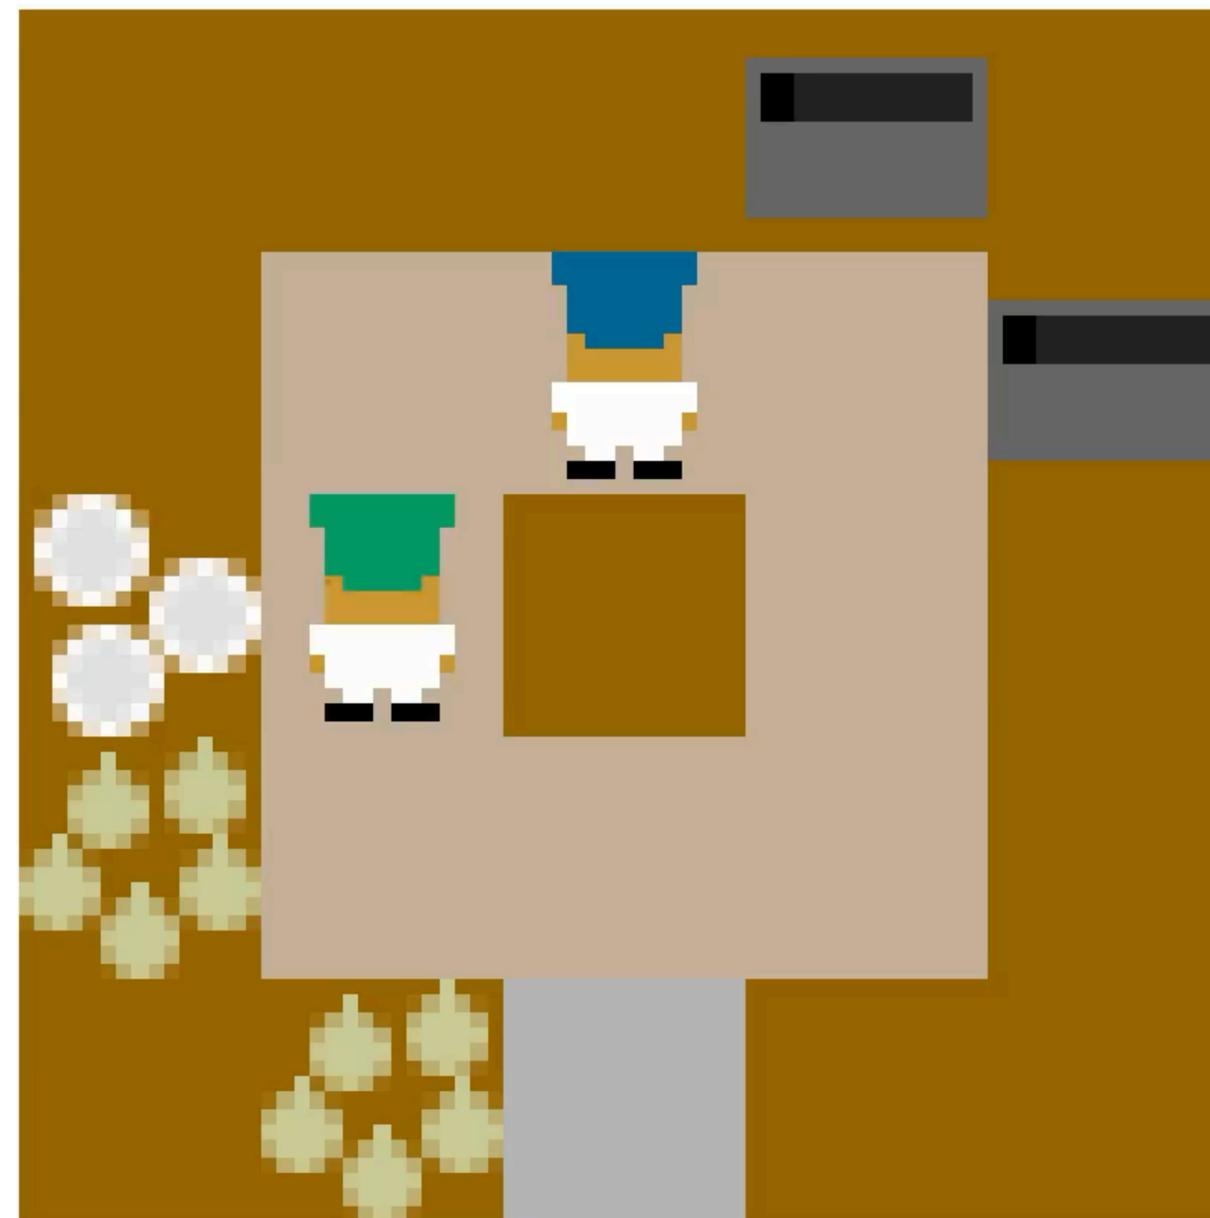
# User Study



# User Study



CoMeDi



Self-Play

CoMeDi enables us to train a **diverse set of conventions**, which can be used to create strong convention-aware agents for human-AI collaboration.

