

LLM-enhanced Agent-Based Model of a Sustainability Game

**Francesco Bertolotti, Federico Carucci,
Luca Mari, Giacomo Buonanno**

*School of industrial engineering, LIUC – Università Cattaneo
Intelligence, Complexity and Technology Lab*

Sabin Roman

*Center for the Study of Existential Risk, Cambridge University
Department of Knowledge Technologies, Jozef Stefan Institute*


NERCCS 2026, 12/03/2026



The Sustainability Game




Game divided into turns, each player plays once per turn



Each player manages an inventory of resources



Different blocks has different features and effects



Each turn players decide what to produce



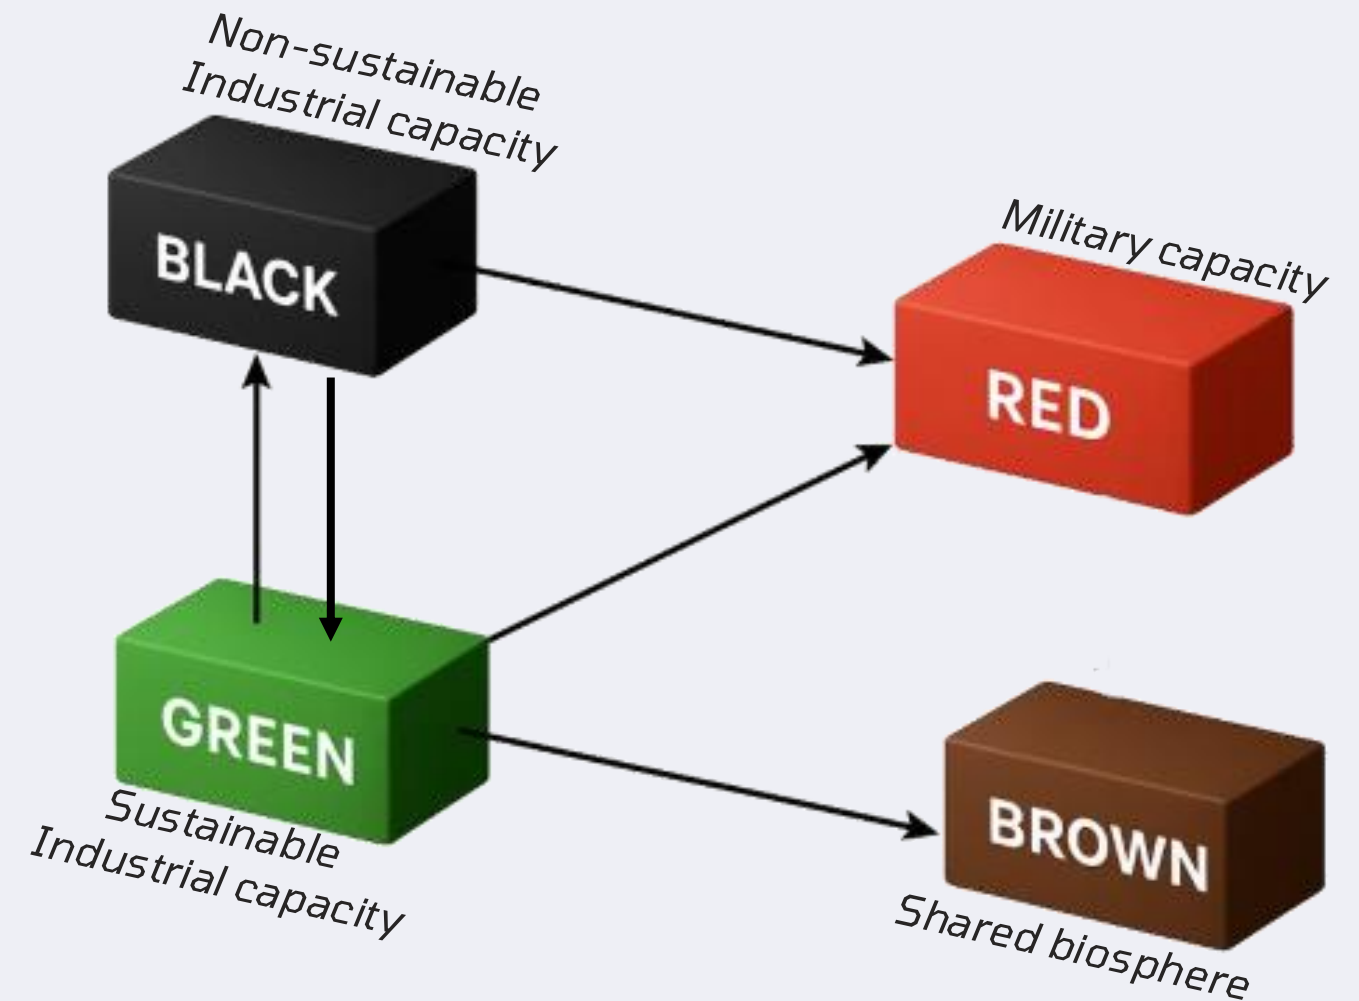
Players are connected to each other through a network



Each turn players decide if and who to attack



Agents pursue either short-term gains or long-term survival strategies





Scenario 1: Extinction

All agents go extinct before reaching the end of the simulation due to resource depletion



Scenario 2: Domination

A single agent militarily defeats all others and emerges as the sole survivor



Scenario 3: Coexistence


Multiple agents successfully survive until the final time step

Overall System Patterns


Research goal



Bertolotti & Roman 2022
Evolutionary dynamics of strategies



Bertolotti & Roman 2024
Optimal combination of strategies and information role



Bertolotti & Roman 2025
*Understand how to use LLM in ABM
Understand how to LLM behaves in complex systems*

LLM-Enhanced Agent

Sensor

Neighbors' block composition Biosphere state

Actuator

Block production Attack targets

Environment

Biosphere state Connected agents

Agent



Blocks
Decisions memory
Biosphere memory

Cognitive process - LLM

System prompt

Rules
Role

Prompt

Status
Strategy
Environment
Forecast

Output

Structured

Decision production prompts

```
if par.strategy:
```

```
    system_prompt += "\n\n YOUR STRATEGY IS: [" + self.strategy + "]"
```

```
if par.prediction_brown:
```

```
    system_prompt += f"\n\n With the current trend of biosphere consumption, in {par.length_forecast} turns, the brown blocks will be {gv.brown_prediction}."
```

```
if par.information_neighbors:
```

```
    system_prompt += f"\n\nThe mean of your neighbors is: black={np.mean([ag.black for ag in self.neighbors]):.1f}, red={np.mean([ag.red for ag in self.neighbors]):.1f}, green={np.mean([ag.green for ag in self.neighbors]):.1f}."
```

```
if par.explanation_production:
```

```
    system_prompt += "\n\n The explanation of your decision can be max 1 sentence."
```

```
content = f"Agent {self.who}: black={self.black}, red={self.red}, green={self.green}\n"
```

```
messages = [{"role": "system", "content": system_prompt},  
            {"role": "user", "content": content}]
```

Agents' strategies



Killer

You produce as many red blocks possible. You ignore prediction about brown decreasing. You attack multiple agents per time, only if you have much more reds then them



Green

You produce only green blocks and very few red blocks. You attack only if you are certain to win. You never create black blocks. You use the greens mostly to create as many browns as possible.



Peace

You produce some red blocks for self defense, just few. You never attack. You make sure you have more green than blacks and green combined. If brown predictions are low, you convert all your blacks into greens and use the greens only to produce browns.



30s'

You first produce as many blacks as you can. Then you produce a lot of reds. You always want to have more reds than the average. If the prediction for the browns gets negative, you stop everything else, convert all the blacks into green and use all the green to produce browns.



Blnc

You produce a balanced number of greens, blacks and reds. You use all the green you have always to create new browns. If the prediction for browns get negative, attack everyone with all you reds and convert your blacks into greens.



Drill

You do not care about the depletion of green blocks. You just want to produce half of your capacity in blacks and the other half in reds. You attack always neighbors weaker than you.



Avgr

You always attack whoever has more blacks than you.

Experimental Design



Parameter Space



Population size



Network connectivity




Forecasting capacity




Implementation



OpenAI API



GPT-4o-mini



Parallelization



Agent Strategies

Exp. 3

killer
green
peace

Exp. 2

thirties
balance

Exp. 1

drillbaby
avenger

Exp. 4

← no strategy




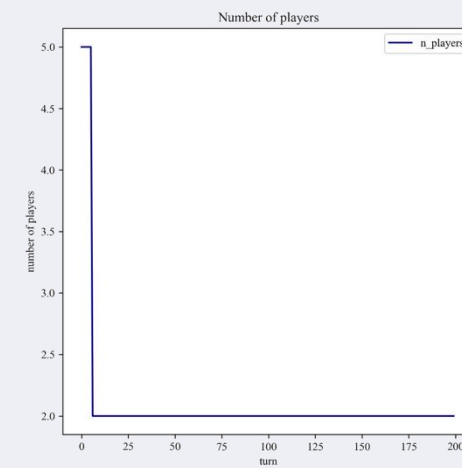
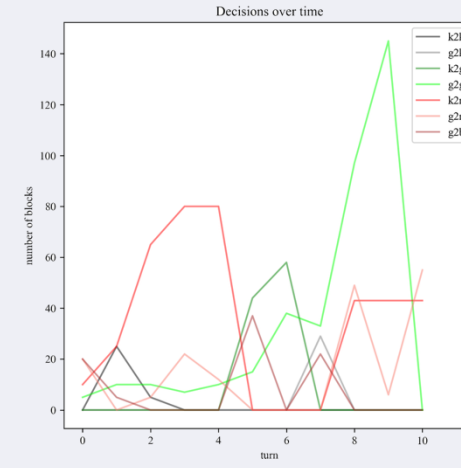



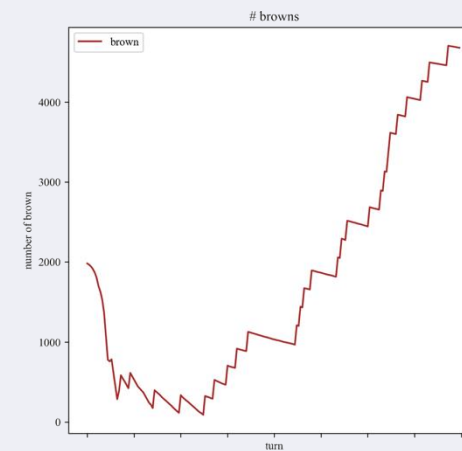
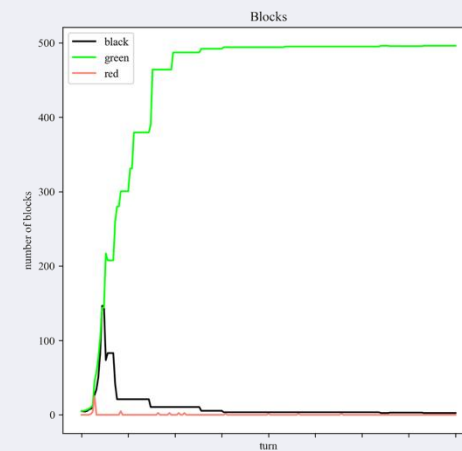

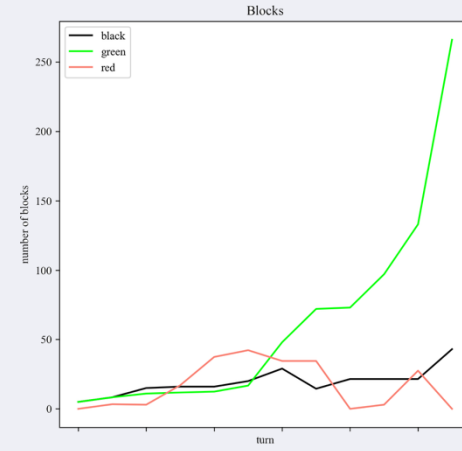
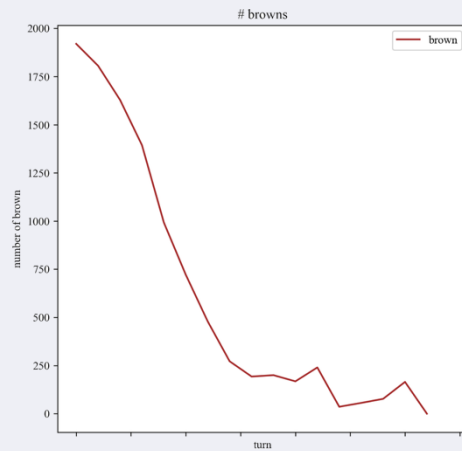
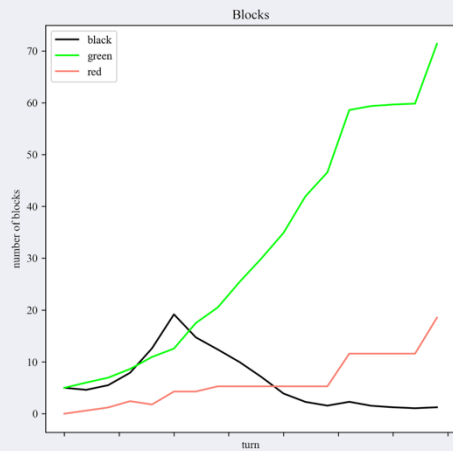
Scenario 1: Extinction



Scenario 2: Domination

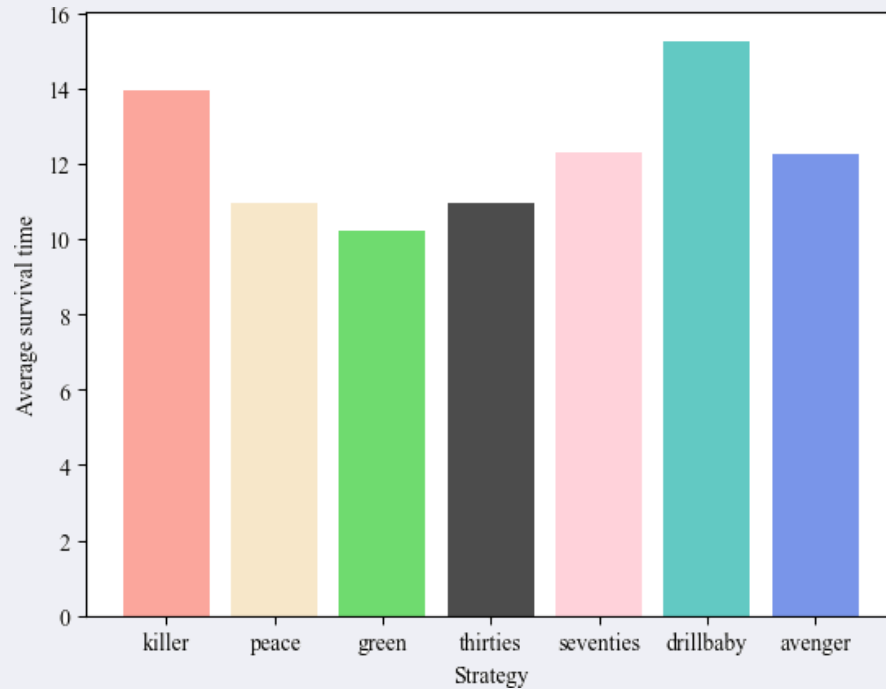


Scenario 3: Coexistence



Results – Aggregate behaviour

Results – experiment 1



Agents are able to coexist only in the **0.5%** of the simulations



The **drillbaby** strategy yielded the **highest** survival time, followed by killer



The **green** strategy resulted in the **lowest** average survival




Aggressive or resource-intensive strategies **outperformed** more **cooperative** ones



Increasing the **number** of players led to a **reduction** in **simulation length**



Larger populations accelerated competition, leading to **quicker** system **destabilization**



Results – experiment 2 and 3

Experiment 2




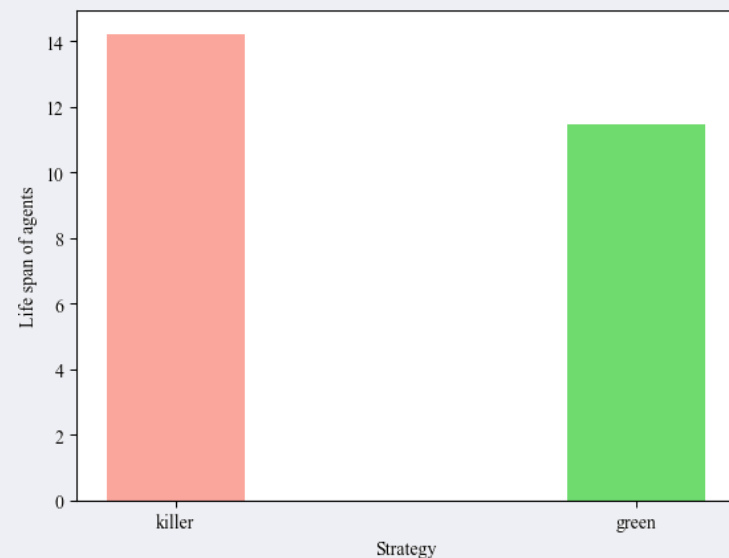
Agents are able to coexist only in the **0.5%** of the simulations



The **killer** strategy yielded the **highest** survival time, 25% higher than green



Mean **average surviving time increases**, but still strong population size effect



Experiment 3




Still **competitions** and attacks



No resources depletion



Few players have been **killed**



Results – experiment 4



Agents are unguided,
without a specific strategy




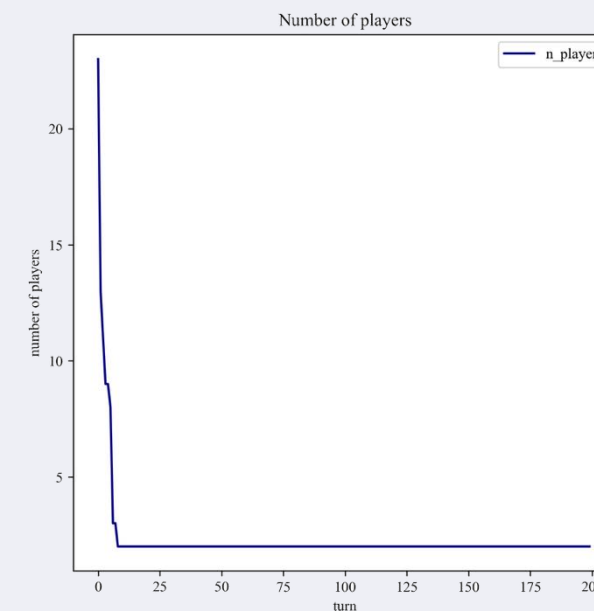

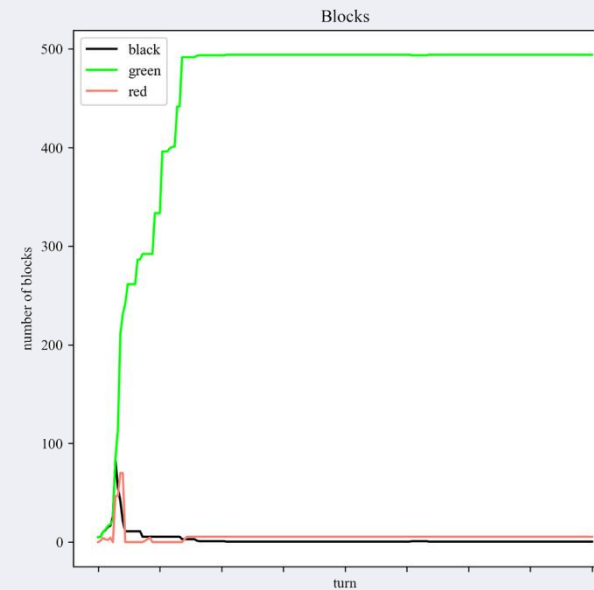
No simulation ended in
domination



38.8% of simulations
ended in coexistence



Extinction often depends on a
single agent's actions



Discussion and Conclusions

LLM-ABM ok...however!

- Computational power
 - Issues with reproducibility
- Dependency to LLM developer
 - Hard generalizability
 - Hard to get understanding
- Need for higher intelligence

Future developments

- Other LLMs
- Other experiments
- Non-trivial strategies
- Findings validation on different settings
- Sensitivity to system prompt



Strategy Influence

- Explicit behavioral guidance alters system balance, with different effects
- Strategies overcome natural behavior, which is more cooperative and sustainable

Population + network

- Higher population decreases chances of survival.
- Less dense networks promote sustainability
- higher link densities are associated with extinction and domination scenarios.

Thank you for your attention!

*LLM-enhanced Agent-Based Model of
a Sustainability Game*

NERCCS 2026, 11/03/2026

**Francesco Bertolotti, Federico Carucci,
Luca Mari, Giacomo Buonanno**

*School of industrial engineering, LIUC – Università Cattaneo
Intelligence, Complexity and Technology Lab*

Sabin Roman

*Center for the Study of Existential Risk, Cambridge University
Department of Knowledge Technologies, Jozef Stefan Institute*

fbertolotti@liuc.it



**Slides are mostly
AI generated!**