For Learning in Symmetric Teams, Local Optima are Global Nash Equilibria

Scott Emmons

International Conference on Machine Learning, 2022

Center for
Human-Compatible
Artificial
Intelligence

BAIR
BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

In Collaboration with

Caspar Oesterheld

Andrew Critch

Vince Conitzer

Stuart Russell

Motivating Example

Common-payoff game

- All players have same payoffs

Symmetric game structure

- Can swap the taxis

Symmetric strategy profile

- Taxis share source code

Types of Symmetry

Symmetric agents & env

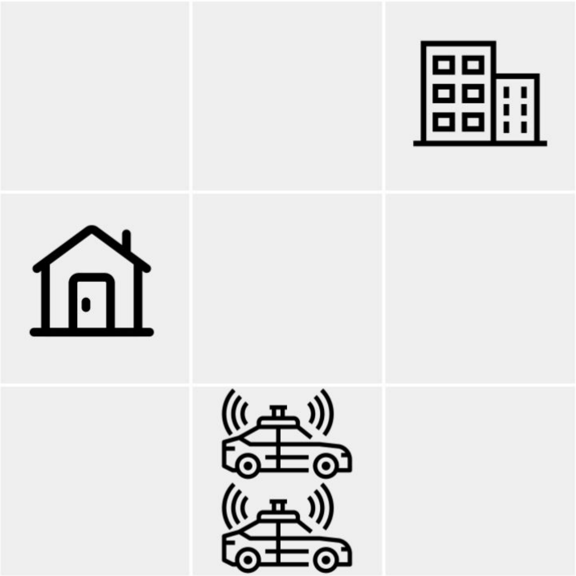

Same initial conditions

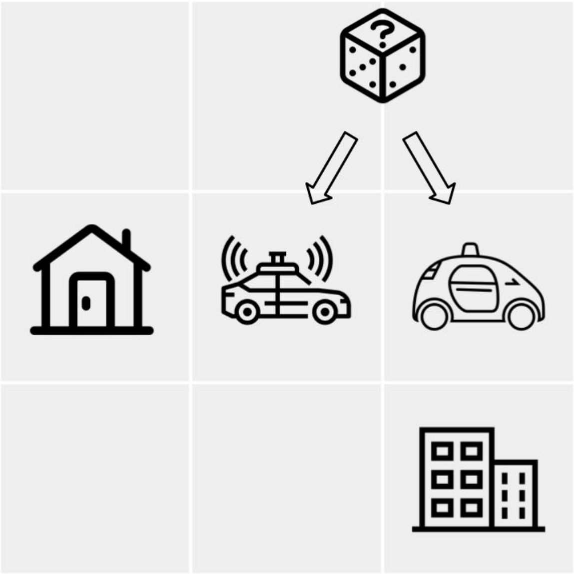

Veil of ignorance

Strategy Profile Landscape

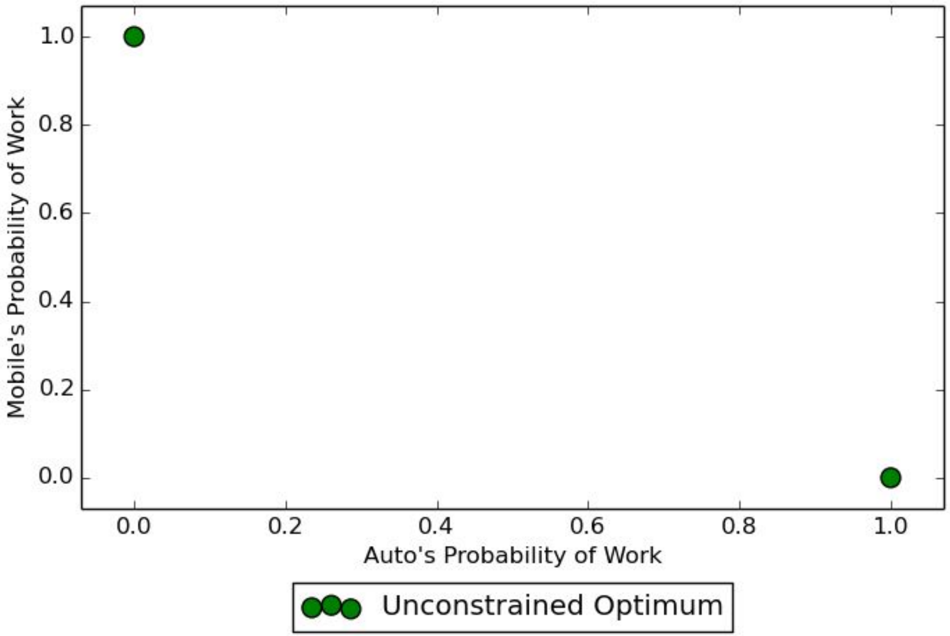

Strategy Profile Landscape

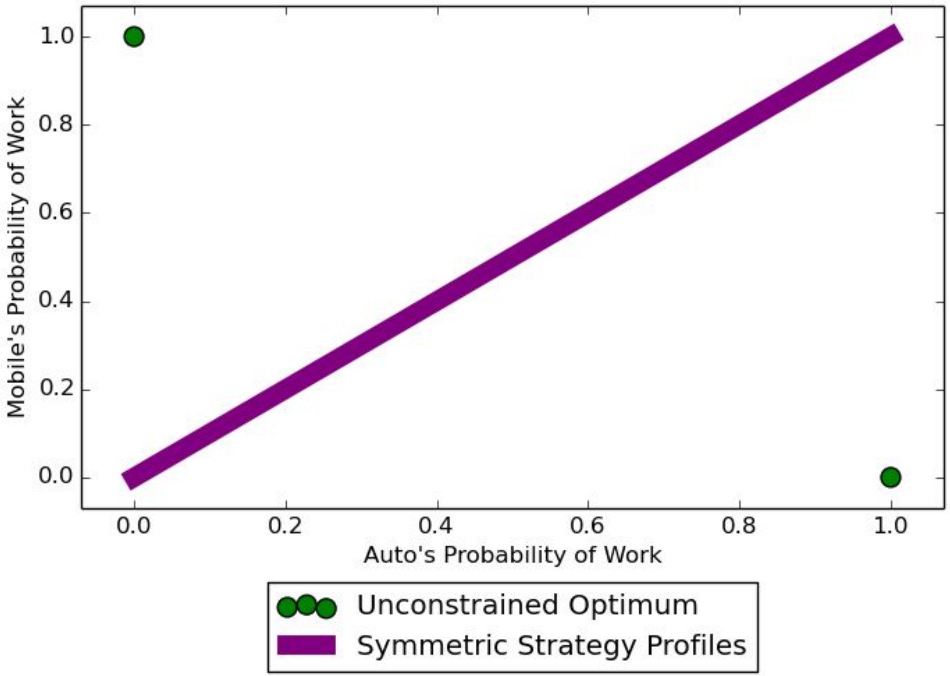

Strategy Profile Landscape

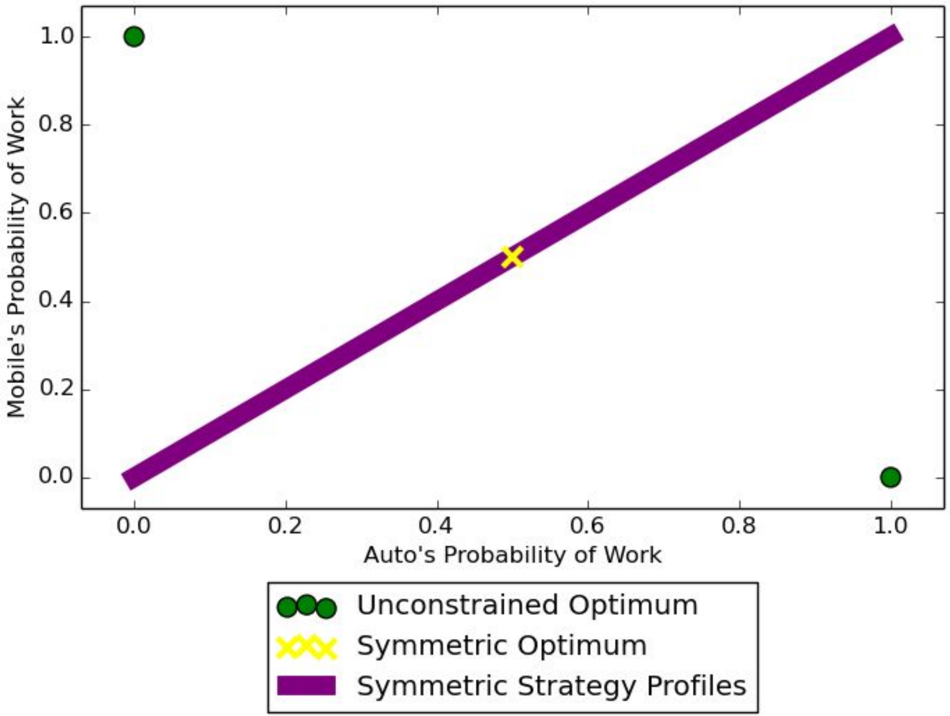

Strategy Profile Landscape

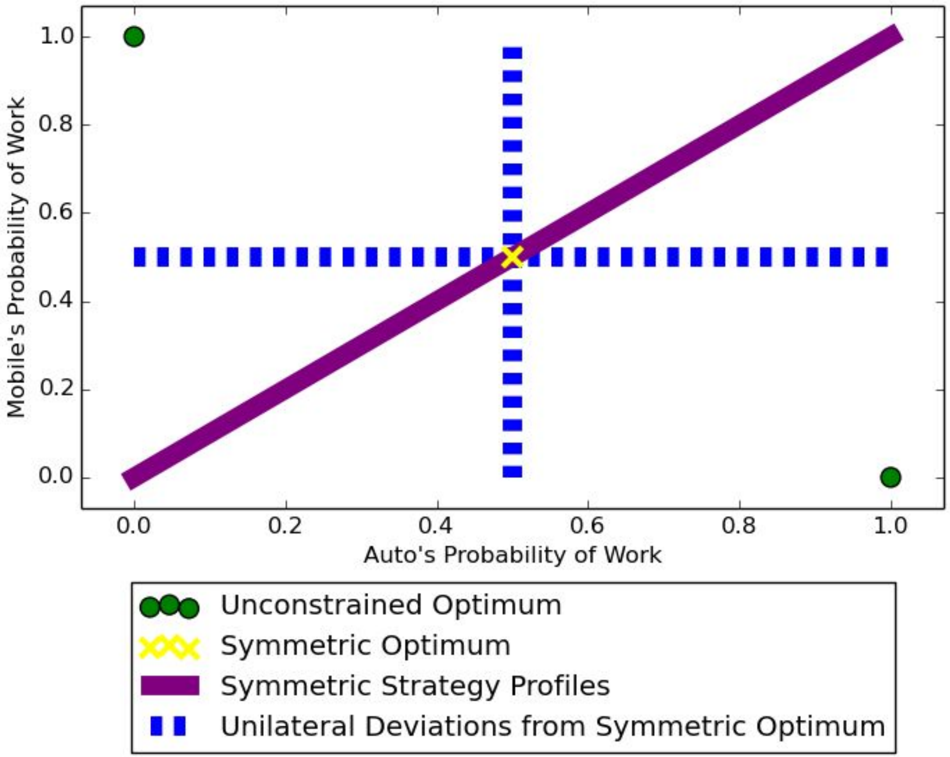

Strategy Profile Landscape

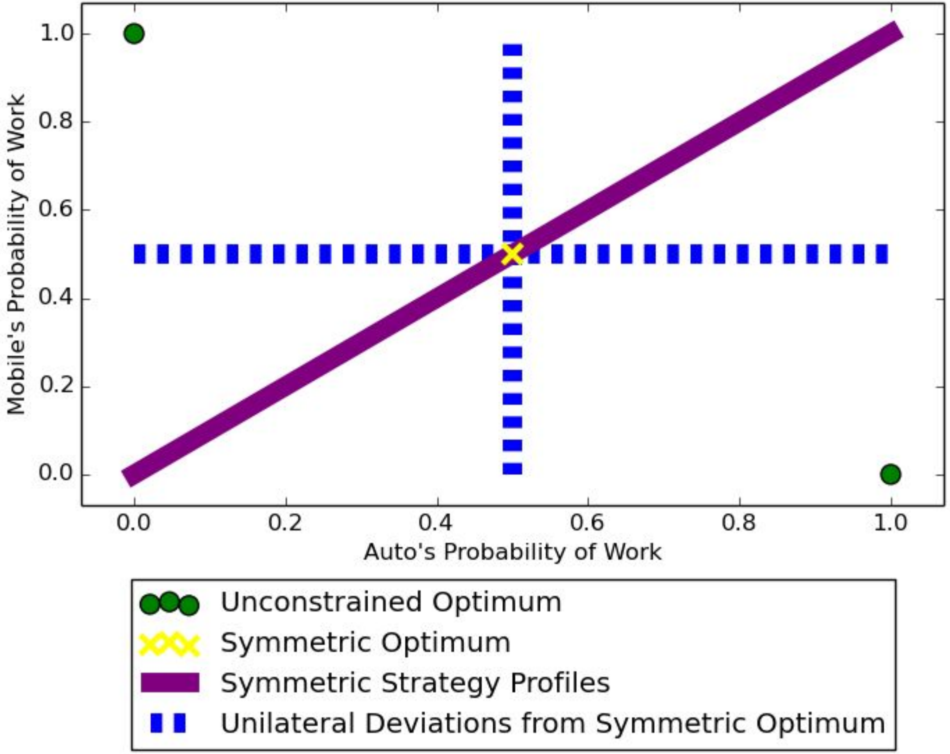

How much can one of *N* players improve the common payoff by unilaterally deviating?

a) No improvement is possible
b) $O(\sqrt{N})$ improvement
c) $O(N)$ improvement
d) $O(N^2)$ improvement

Strategy Profile Landscape

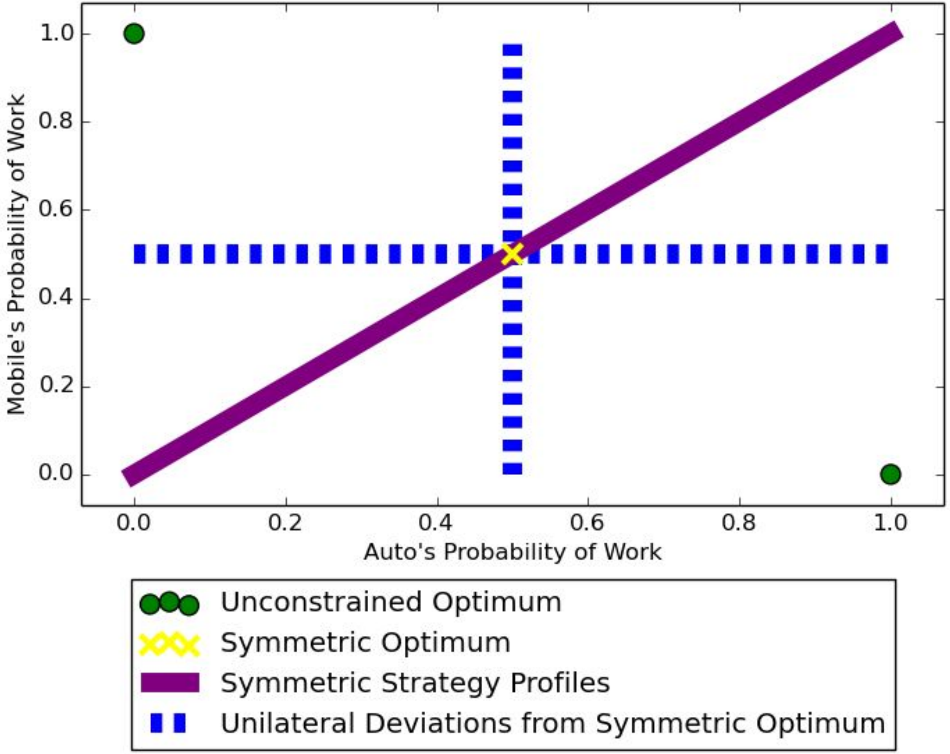

How much can one of *N* players improve the common payoff by unilaterally deviating?

a) **No improvement is possible**

~~b) $O(\sqrt{N})$ improvement~~

~~c) $O(N)$ improvement~~

~~d) $O(N^2)$ improvement~~

Strategy Profile Landscape

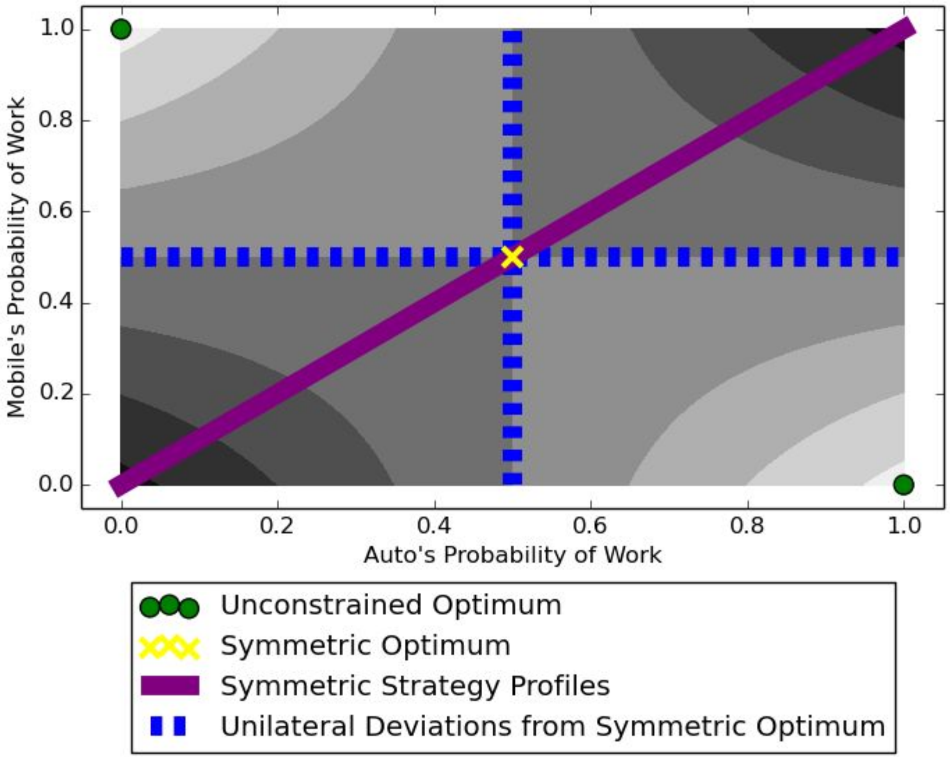

	Home	Work
Home	1	2
Work	2	1

Local $\rightarrow$ Global Guarantee

In Plain English

Locally optimal symmetric collaboration is a *global* Nash equilibrium.

Theorem

Let $\mathcal{G}$ be a normal-form (or extensive-form) game with common payoff.
Then any *locally* optimal $\mathcal{P}$-invariant strategy profile is a *global* Nash equilibrium.

Note: $\mathcal{P}$-invariance is a group-theoretic notion of symmetry generalizing:

- Anonymous games
- Transitive games
- Totally symmetric games

Implications of Theorem

Team Theory

No individual can innovate from a symmetric local optimum to improve the common payoff.

Multi-agent Reinforcement Learning

Iterated best response can’t improve a local symmetric optimum.

(Adversarial) Team Games

Result extends to arbitrary number of teams.

Robustness of Result

The result degrades smoothly, giving an ε-Nash equilibrium for:

Approximate Solutions

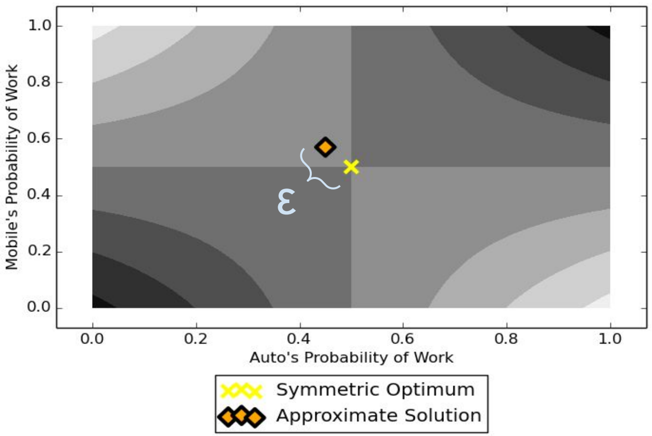

Payoff Perturbations

	Home	Work
Home	1+ε, 1-ε	2-ε, 2-ε
Work	2+ε, 2+ε	1-ε, 1+ε

(In)stability to Joint (Possibly-Asymmetric) Deviation

<u>Theorem</u> (for non-degenerate games)

A local symmetric optimum is locally optimal among possibly-asymmetric strategy profiles *if and only if it is deterministic.*

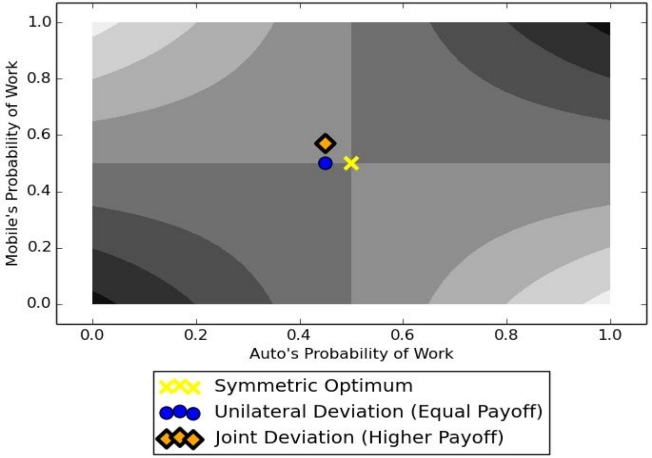

Experimentally, up to 60% are mixed, i.e., *unstable*!

Conclusion

We give conditions for *stability* and *instability*

<u>Future Work</u>

- Behavioral strategies
- Continuous actions & function approximation
- Sequential decision making benchmarks