Learnable Mixed Nash Equilibria Are Collectively Rational

Geelon So and Yian Ma

1. In **multi-agent systems**, the agents must often **learn from experience** how to behave; they don't innately know 'how the game works'.

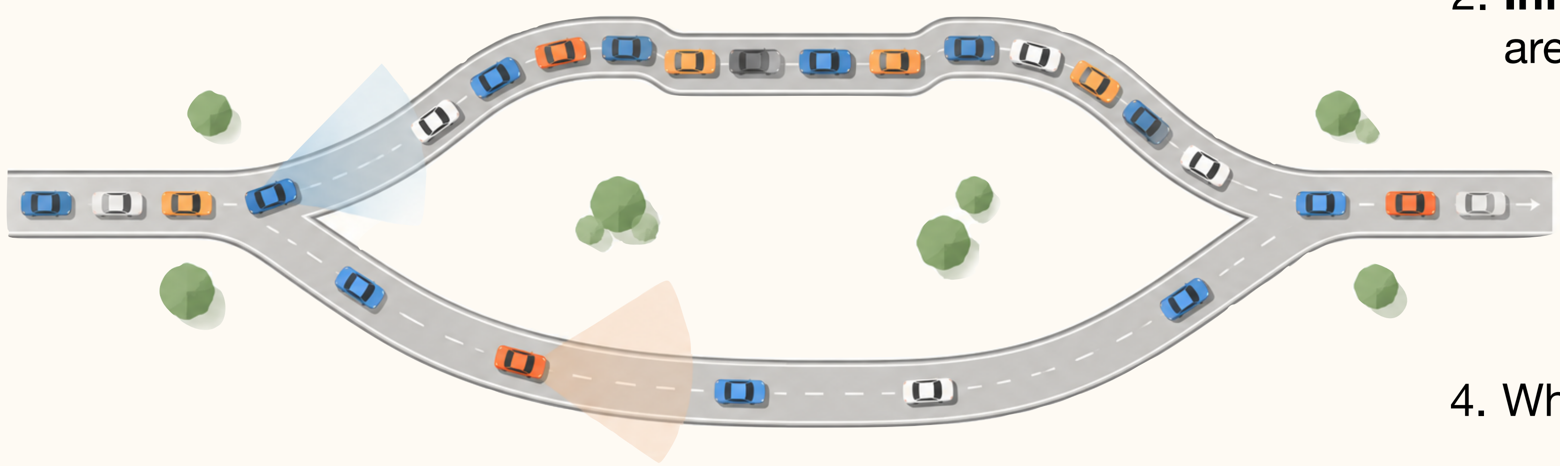

2. **Information asymmetry** and **mis-** or **unaligned interests** are two fundamental challenges inherent to games.

3. Actions can seem optimal from the **individual** standpoint, but from a **collective** one, they do not need to be rational.

4. Which **solutions** can the agents actually discover?

1. LEARNING IN GAMES

General-sum, N-player normal-form games:

- Each agent $n \in [N]$ has mixed strategy space Ω_n.
- The joint strategy space is $\mathbf{\Omega} = \Omega_1 \times \cdots \times \Omega_N$.
- The utility of the nth agent is $f_n : \mathbf{\Omega} \to \mathbb{R}$.

Agents have limited ability to cooperate:

- Each agent only knows its own utility.
- It doesn't know *what others want* nor *how they learn*.
- Agents cannot communicate.

Learning dynamics:

For $t = 1, 2, \ldots$:

- Each agent plays an action $x_n(t) \in \Omega_n$.
- The joint action $\mathbf{x}_n(t) \in \mathbf{\Omega}$ is revealed to everyone.

Example (preconditioned gradient ascent-ascent)**:**

$$x_n(t+1) = x_n(t) + \eta H_n^{-1} \nabla_n f_n\big(\mathbf{x}(t)\big) \,.$$

2. UNCOUPLED DYNAMICS

Theorem (Hart and Mas-Colell 2003)**.** No uncoupled dynamics *robustly* converges to all Nash equilibria.

- *There is a family of games such that, for any learning dynamics, there is a game for which the dynamics almost surely avoids the unique Nash equilibrium.*

Linear stability analysis: Standard game dynamics have Jacobians related to the following matrix

$$\mathbf{H}^{-1}\mathbf{J},$$

where $\mathbf{J}_{nm} = \nabla^2_{nm} f_n$, $\mathbf{H} = \mathbf{blk\text{-}diag}(H_1, \ldots, H_N)$, and $H_n \succ 0$.

3. RATIONALITY

Individual rationality: A solution $\mathbf{x}^* \in \mathbf{\Omega}$ is a *Nash equilibrium* if fixing everyone else's actions

$$x_n^* \in \arg\max_{x_n \in \Omega_n} f_n(x_n; \mathbf{x}^*_{-n}) \,.$$

Collective rationality: A solution $\mathbf{x}^*$ is *weakly Pareto optimal* if there is no way to strictly improve everyone.

4. CONVERGENCE

Dynamical stability: The game is *uniformly stable* if all eigenvalues of the Jacobian are imaginary

$$\forall \mathbf{H}, \qquad \mathrm{spec}\big(\mathbf{H}^{-1}\mathbf{J}\big) \subset i\mathbb{R} \,.$$

Theorem. Let $\mathbf{x}^*$ be a mixed Nash equilibrium. If the game is uniformly stable, then the equilibrium is weakly Pareto optimal (up to strategic equivalence).

- *While Hart and Mas-Colell (2003) showed that we cannot converge to every Nash equilibrium, we show that the ones we can always learn are socially efficient.*

Theorem. Under *incremental smoothed best-response dynamics*, non-degenerate and uniformly stable Nash equilibria admit convergence at a rate of

$$O(T^{-1/2}).$$

REFERENCES

Hart, S. and Mas-Colell, H. *Uncoupled Dynamics Do Not Lead to Nash Equilibrium.* The American Economic Review (2003).