Strategically Robust Game Theory via Optimal Transport

Nicolas Lanzetti

joint with S. Fricker, S. Bolognani, F. Dörfler, and D. Paccagnan

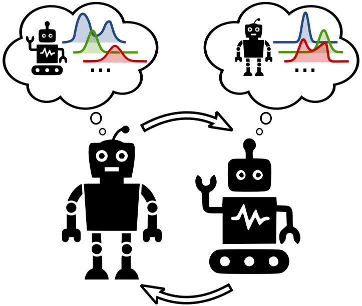

ETH zürich Caltech Imperial College London

the main source of uncertainty is the behavior of the other agents
but the other agents are strategic too

autonomous driving

stock market

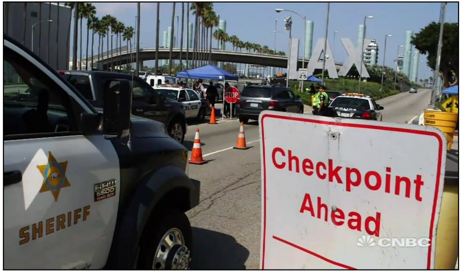

security

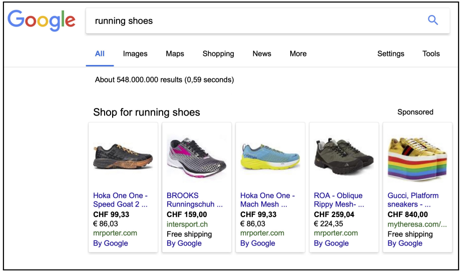

online bidding

limited information

limited rationality

limited computation

limited sensing

if you can, model them...

▷ limited information ⟹ Bayesian GT

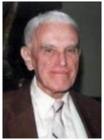

John Harsanyi


GAMES WITH INCOMPLETE INFORMATION PLAYED BY "BAYESIAN" PLAYERS, I–III

Part I. The Basic Model

JOHN C. HARSANYI

1966

▷ limited rationality ⟹ behavioural GT

Maurice Allais

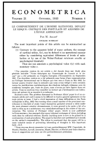
ECONOMETRICA

LE COMPORTEMENT DE L'HOMME RATIONNEL DEVANT LE RISQUE: CRITIQUE DES POSTULATS ET AXIOMES DE L'ECOLE AMERICAINE

1953

▷ limited computation ⟹ algorithmic GT

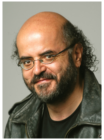

Papadimitriou

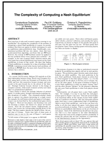
The Complexity of Computing a Nash Equilibrium

2008

. . . but you almost never can!

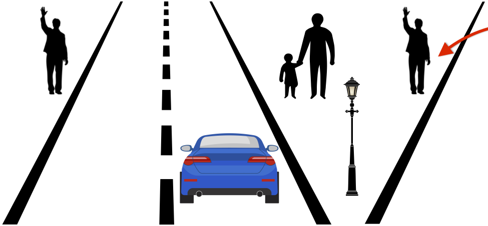

	Wait	Cross
Maintain	$(10, -1)$	$(-50, -100)$
Decel	$(9, -1)$	$(-5, -10)$
Stop	$(0, -1)$	$(0, 10)$

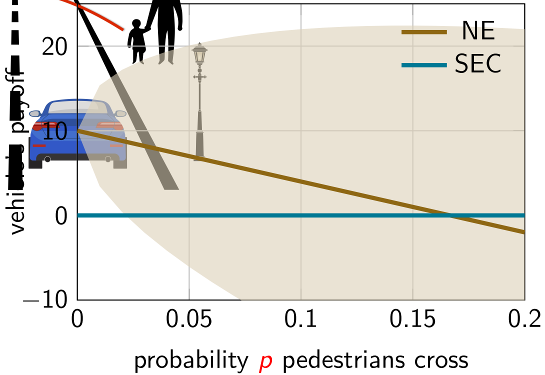

Nash is not robust, security strategy is too robust

can we break this barrier?

yes, with strategically robust equilibria

- ▷ players $i = 1, \dots, N$
- ▷ action space $\mathcal{A} = \mathcal{A}^1 \times \cdots \times \mathcal{A}^N$
- ▷ strategy space $\Delta^i = \mathcal{P}(\mathcal{A}^i)$
- ▷ expected utility $U^i(p^i, p^{-i}) = \mathbb{E}^{p^i, p^{-i}}[u^i(a^i, a^{-i})]$
- ▷ ambiguity sets $B^i_\varepsilon(p^{-i}) \subseteq \Delta^{-i}$

p^{-i}

$B^i_\varepsilon(p^{-i})$

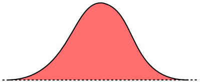

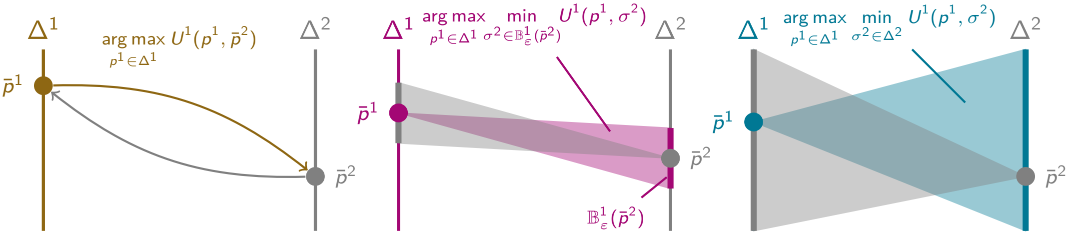

Nash equilibria
i.e., SRE with $\varepsilon = 0$

strategically robust equilibria
(SRE)

security strategies
i.e., SRE with $\varepsilon = \infty$

Definition (Strategically robust equilibrium – SRE):

$$(\bar{p}^1, \ldots, \bar{p}^N) \quad \text{s.t.} \quad \bar{p}^i \in \arg\max_{p^i \in \Delta^i} \min_{\sigma^{-i} \in B^i_\varepsilon(\bar{p}^{-i})} U^i(p^i, \sigma^{-i})$$

Challenges:

1) constrained worst-case optimization over mixed strategies (probabilities over actions)

2) fixed-point of optimization problems over mixed strategies

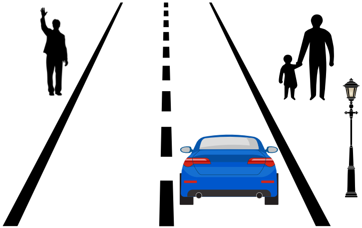

	WAIT	CROSS
MAINTAIN	$(10, -1)$	$(-50, -100)$
DECEL	$(9, -1)$	$(-5, -10)$
STOP	$(0, -1)$	$(0, 10)$

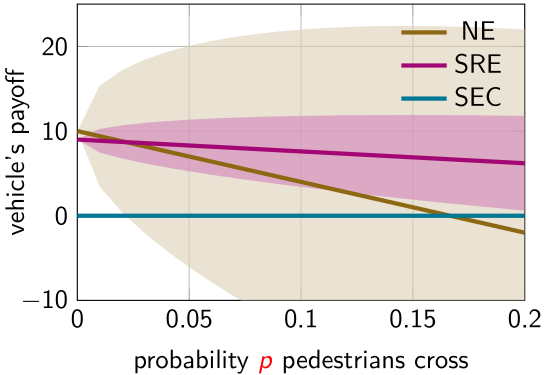

ambiguity sets

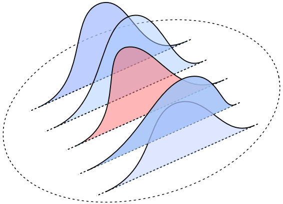

desiderata:

- *expressivity:* protect against different distributions
- *tractability:* give rise to an "easy" problem
- *interpolate:* $B_0^i(p^{-i}) = p^{-i}$, $\lim_{\varepsilon\to\infty} B_\varepsilon^i(p^{-i}) = \Delta^{-i}$

possible approaches: optimal transport, KL-divergence, moment based, ...

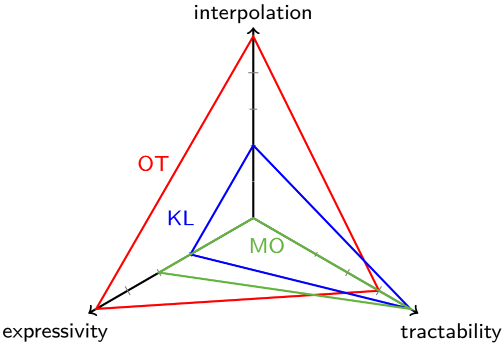

$$B^i_\varepsilon(p^{-i}) = \{\sigma^{-i} \in \mathcal{P}(\mathcal{A}^{-i}) \ \text{ s.t. } \ W(p^{-i}, \sigma^{-i}) \leq \varepsilon\}$$

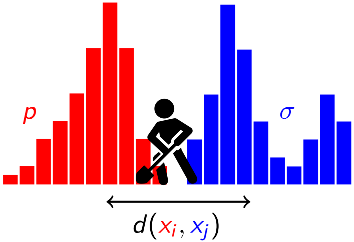

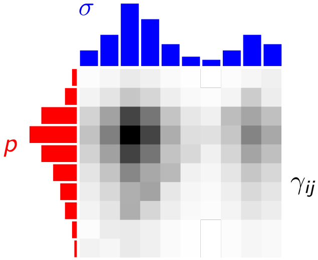

$$W(p, \sigma) = \min_{\gamma \in \mathbb{R}^{|\mathcal{A}| \times |\mathcal{A}|}_{\geq 0}} \sum \sum d(x_i, x_j)\gamma_{ij}$$

$$\text{s.t.} \quad \sum_i \gamma_{ij} = p_j, \quad \sum_j \gamma_{ij} = \sigma_i$$

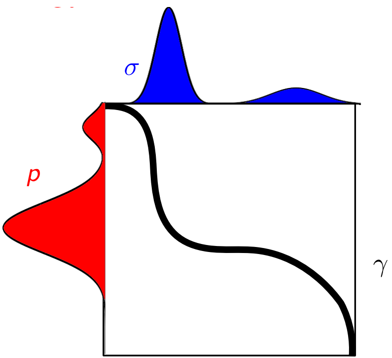

$$W(p,\sigma) = \inf_{\gamma} \int_{\mathcal{X}\times\mathcal{X}} d(x,y)\mathrm{d}\gamma(x,y)$$

$$\text{s.t.} \quad \gamma \text{ has first marginal } p$$

$$\gamma \text{ has second marginal } \sigma$$

existence

recall SRE definition:

Definition (strategically robust equilibrium – SRE):

$$(\bar{p}^1, \dots, \bar{p}^N) \quad \text{s.t.} \quad \bar{p}^i \in \arg\max_{p^i \in \Delta^i} \min_{\sigma^{-i} \in B^i_\varepsilon(\bar{p}^{-i})} U^i(p^i, \sigma^{-i})$$

We will use *either* of the following assumptions:

A1: The action spaces $\{\mathcal{A}^i\}_{i=1}^N$ are finite;

A2: Each action space $\mathcal{A}^i$ is a compact subset of $\mathbb{R}^n$
Each payoff function u^i is continuous.

Theorem (existence): Assume A1 or A2. For any $\varepsilon \geq 0$, a strategically robust equilibrium exists.

same assumptions needed for existence mixed Nash!

Challenges: coupled DRO problems, Nash's proof does not extend

finite games: computation

long line of work: Chen, Daskalakis, Goldberg, Papadimitriou, Savani, ...

culminated with: mixed Nash equilibria are PPAD-complete, for $N \geq 2$

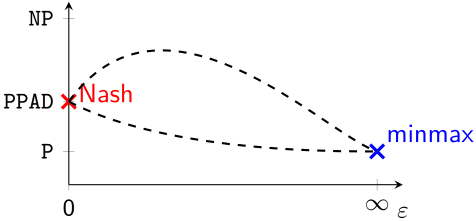

Theorem (computation, part 1): For any $\varepsilon \geq 0$ and $N \geq 2$ the computational complexity of SRE is in PPAD.

no harder than mixed Nash... in fact, even easier – stay tuned!

Theorem (computation, part 2): For any $\varepsilon \geq 0$, SRE are found solving multilinear complementarity problem.

recall that multilinear CP asks for $x \in \mathbb{R}^n$ such that

$$0 \leq x \perp F(x) \geq 0 \quad \Leftrightarrow \quad x_i F_i(x) = 0$$

- ▷ just like for mixed Nash!
  - *linear* for 2 player games
  - *linear* for special classes, e.g., polymatrix
- ▷ use off-the-shelf solvers (e.g., PATH solver)
- ▷ this is what we used for all numerics to follow

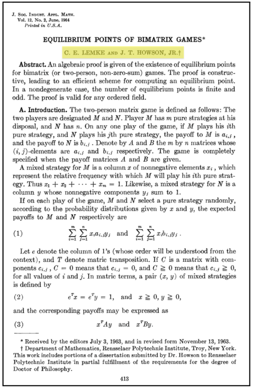

J. Soc. Indust. Appl. Math.
Vol. 12, No. 2, June, 1964
Printed in U.S.A.

EQUILIBRIUM POINTS OF BIMATRIX GAMES*

C. E. LEMKE and J. T. HOWSON, JR.†

Abstract. An algebraic proof is given of the existence of equilibrium points for bimatrix (or two-person, non-zero-sum) games. The proof is constructive, leading to an efficient scheme for computing an equilibrium point. In a nondegenerate case, the number of equilibrium points is finite and odd. The proof is valid for any ordered field.

A. Introduction. The two-person matrix game is defined as follows: The two players are designated M and N. Player M has m pure strategies at his disposal, and N has n. On any one play of the game, if M plays his ith pure strategy, and N plays his jth pure strategy, the payoff to M is $a_{i,j}$, and the payoff to N is $b_{i,j}$. Denote by A and B the m by n matrices whose (i, j)-elements are $a_{i,j}$ and $b_{i,j}$ respectively. The game is completely specified when the payoff matrices A and B are given.

A mixed strategy for M is a column x of nonnegative elements x_i, which represent the relative frequency with which M will play his ith pure strategy. Thus $x_1 + x_2 + \cdots + x_m = 1$. Likewise, a mixed strategy for N is a column y whose nonnegative components y_j sum to 1.

If on each play of the game, M and N select a pure strategy randomly, according to the probability distributions given by x and y, the expected payoffs to M and N respectively are

$$(1) \qquad \sum_{i=1}^{m} \sum_{j=1}^{n} x_i a_{i,j} y_j \quad \text{and} \quad \sum_{i=1}^{m} \sum_{j=1}^{n} x_i b_{i,j} y_j .$$

Let e denote the column of 1's (whose order will be understood from the context), and T denote matric transposition. If C is a matrix with components $c_{i,j}$, $C = 0$ means that $c_{i,j} = 0$, and $C \geqq 0$ means that $c_{i,j} \geqq 0$, for all values of i and j. In matric terms, a pair (x, y) of mixed strategies is defined by

$$(2) \qquad e^T x = e^T y = 1, \quad \text{and} \quad x \geqq 0, y \geqq 0,$$

and the corresponding payoffs may be expressed as

$$(3) \qquad x^T A y \quad \text{and} \quad x^T B y.$$

* Received by the editors July 3, 1963, and in revised form November 13, 1963.

† Department of Mathematics, Rensselaer Polytechnic Institute, Troy, New York.
This work includes portions of a dissertation submitted by Dr. Howson to Rensselaer Polytechnic Institute in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

413

numerics

SRE avoids Braess' Paradox

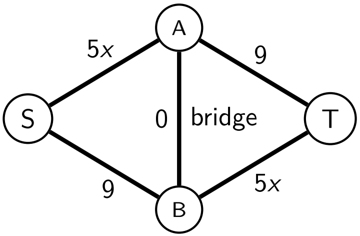

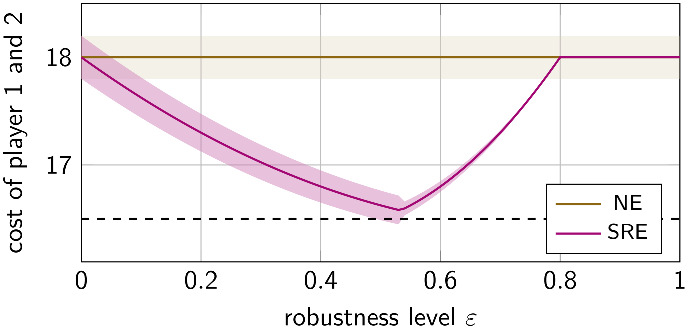

	cooperate	not cooperate
cooperate	$(0.6, 0.6)$	$(0.6, 1)$
not cooperate	$(1, 0.6)$	$(0, 0)$

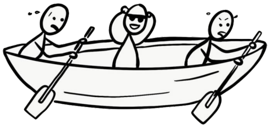

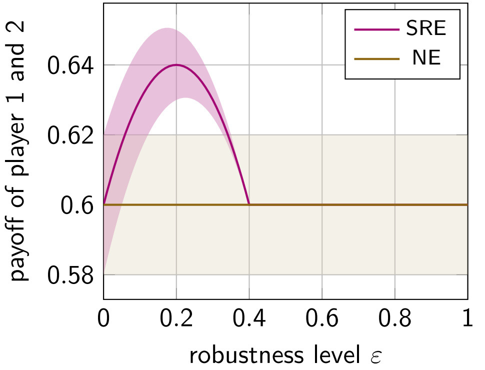

	inspect	not inspect
shirk	$(0, -5)$	$(10, -10)$
work	$(5, 0)$	$(5, 5)$

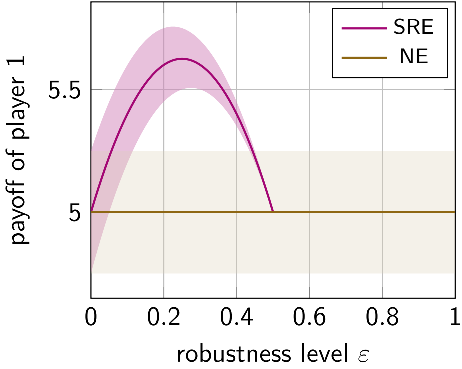

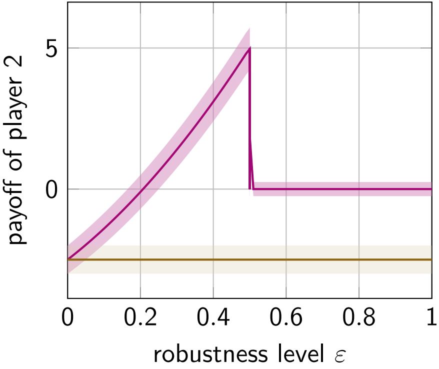

continuous games


Challenges:

- ▷ distributions over continuous actions
  - ⚠ the space of mixed strategies is infinite-dimensional
- ▷ even evaluating $W(p, \sigma)$ is difficult
  - ⚠ the Wasserstein distance itself is an infinite-dimensional LP

Ideas:

- ▷ we prove the existence on pure strategically robust equilibria
- ▷ we use duality theory for distributionally robust optimization to compute equilibria efficiently

⚠ the space of mixed strategies is infinite-dimensional

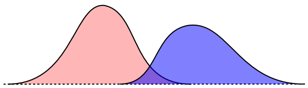

Consider concave games:

- ▷ compact and convex action spaces
- ▷ the payoffs $(a^i, a^{-i}) \mapsto u(a^i, a^{-i})$ are continuous
- ▷ the payoffs $a^i \mapsto u(a^i, a^{-i})$ are concave for fixed a^{-i}

Theorem (pure SRE): Consider a concave game and a robustness level ε. Then, there is a *pure* strategically robust equilibrium with robustness level ε.

thus, we need to look for a finite-dimensional object

⚠ evaluating the Wasserstein distance is difficult

For a concave game $\mathcal{G}$, consider a surrogate concave game $\mathcal{G}_\varepsilon$ with:

▷ the augmented action space $\mathcal{A}^i \times [0, M]$ (M is large constant)

▷ the modified payoffs

$$\tilde{u}_\varepsilon((a^i, \lambda^i), (a^{-i}, \lambda^{-i})) = \min_{\hat{a}^{-i} \in \mathcal{A}^{-i}} \{u^i(a^i, \hat{a}^{-i}) - \lambda^i d(a^{-i}, \hat{a}^{-i})\} - \lambda^i \varepsilon$$

▷ $\mathcal{G}_\varepsilon$ is itself a concave game

Theorem (computation):

$$(\bar{a}_1, \ldots, \bar{a}_N) \text{ is pure SRE of } \mathcal{G} \text{ with robustness level } \varepsilon$$
$$\Longleftrightarrow$$
$$((\bar{a}_1, \bar{\lambda}_1), \ldots, (\bar{a}_N, \bar{\lambda}_N)) \text{ is a pure NE of } \mathcal{G}_\varepsilon \text{ for some } (\bar{\lambda}_1, \ldots, \bar{\lambda}_N)$$

use tools “à la Nash” to compute strategically robust equilibria

(and duality to solve the minimum efficiently)

N firms compete in T markets

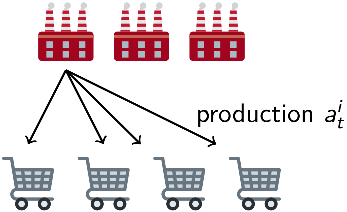

prod. cost of firm i

$$\text{profit}^i(a^i, a^{-i}) = \sum_{t=1}^{T} a_t^i \left(\alpha_t - \beta_t \sum_{i=1}^{N} a_t^i\right) - c^i a_t^i$$

price in market t

production constraints: $\sum_t a_t^i \in [0, K^i]$

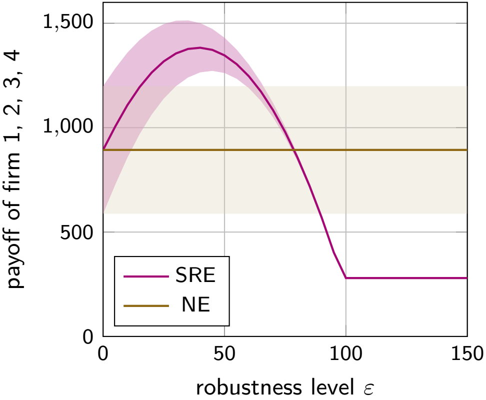

same trend if firms are asymmetric

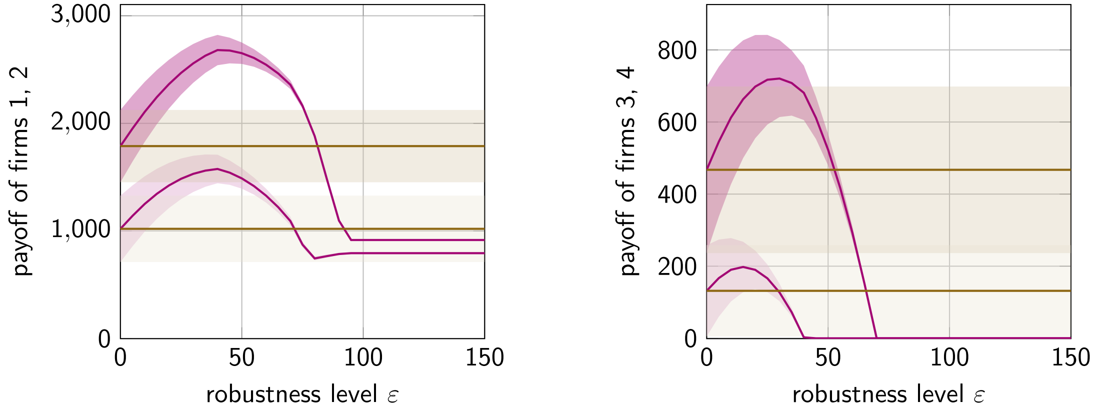

SRGT in the wild: training of collaborative AI agents

joint work with Chengrui Qu, Yizhou Zhang, and Eric Mazumdar state-of-the-art for multi-agent AI training: each agent runs their own policy optimization (I-PPO)

SRPO: each agent runs a risk-averse PPO, assuming they play against a fictitious opponent

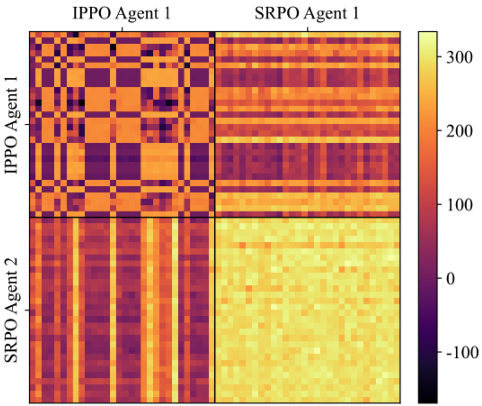

Qu, Zhang, Lanzetti, Mazumdar, *Training Generalizable Collaborative Agents via Strategic Risk Aversion*, arxiv preprint, 2026

Tag

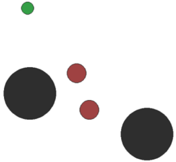

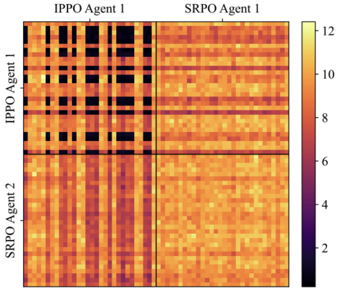

Hanabi

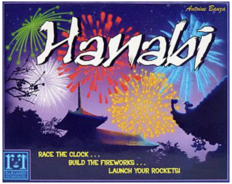

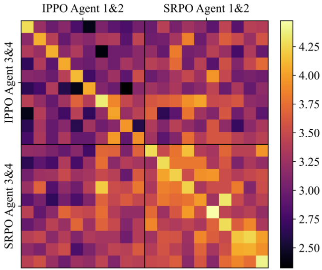

Multi-LLM-Agent debate

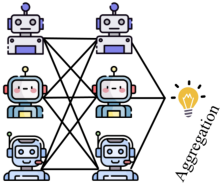

Method	Q0.5B	Q0.6B	Q3B	Q4B
SRPO	**0.405**	**0.671**	**0.632**	**0.917**
IPPO	0.378	0.587	0.552	0.901
Improvement (%)	+7.14%	+14.31%	+14.49%	+1.78%

Qu, Zhang, Lanzetti, Mazumdar, *Training Generalizable Collaborative Agents via Strategic Risk Aversion*, arxiv preprint, 2026

SRGT in the wild:

electric vehicle charging

joint work with Andreas Feik, Saverio Bolognani, Florian Dörlfer, and Dario Paccagnan agents' goal: charge their EVs by θ^i while minimizing electricity costs:

$$J^i\left(a^i, \frac{1}{N}\sum_{i=1}^{N} a_k^i\right) = \sum_{k=1}^{n} a_k^i \text{price}_k\left(\frac{1}{N}\sum_{i=1}^{N} a_k^i\right) \qquad \text{price}_k(\sigma) = \frac{\sigma + d_k}{\kappa_k},$$

$$\mathcal{X}^i = \left\{ a^i \in \mathbb{R}^n : \begin{array}{ll} 0 \le a_k^i \le \bar{a}_k^i, & \forall k = 1, \dots, n \\ \sum_{k=1}^{n} a_k^i \ge \theta^i & \end{array} \right\}.$$

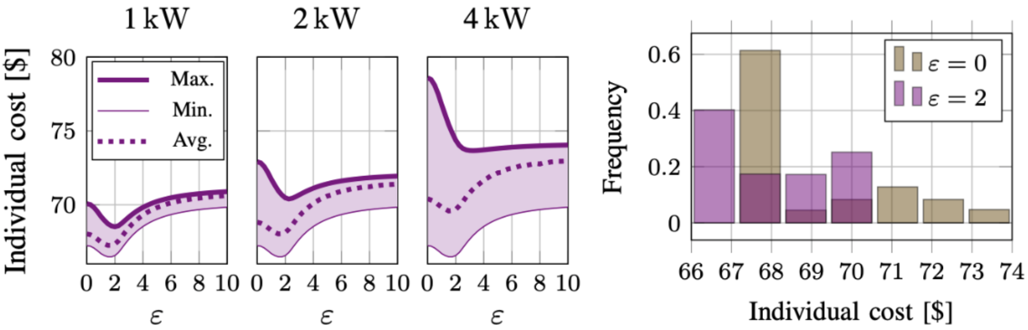

SRGT in the wild:

robust trajectory generation for autonomous vehicles

joint work with Victor Qin

agents' goal: reach their destination while minimizing travel time and avoiding collisions

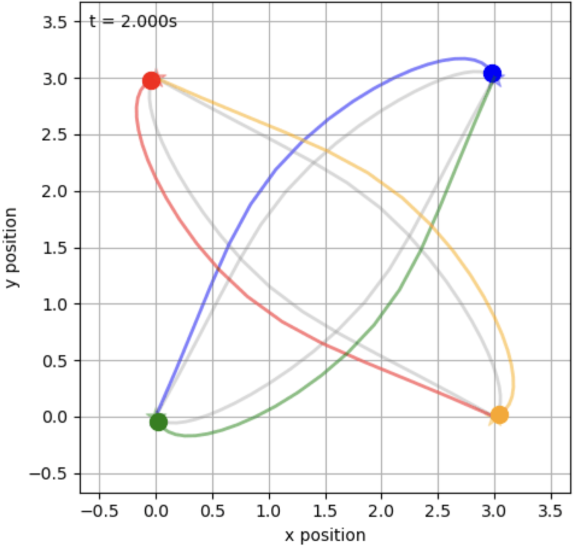

we developed strategically robust equilibria that. . .

. . . protect agents against deviations in the other's behavior

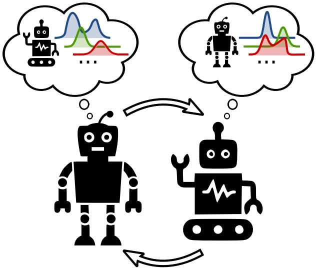

. . . often yield higher payoff for all agents

. . . come at no additional cost compared to NE, despite the extra robustness