

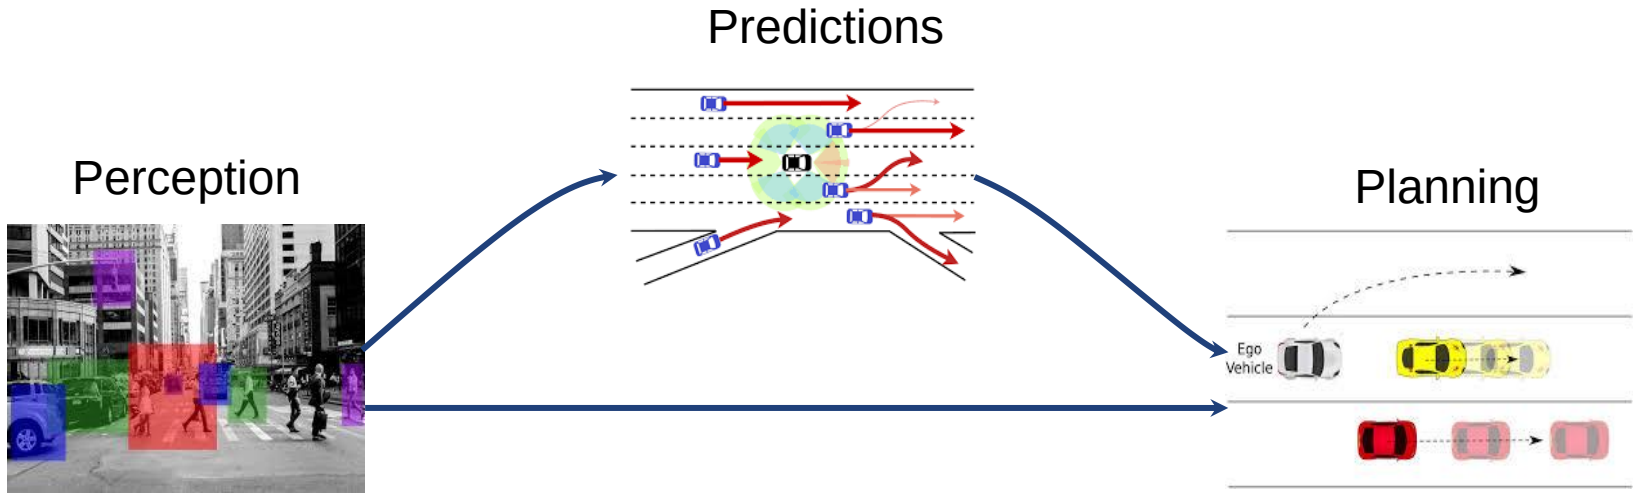
Structure and Solution of Urban Driving Games

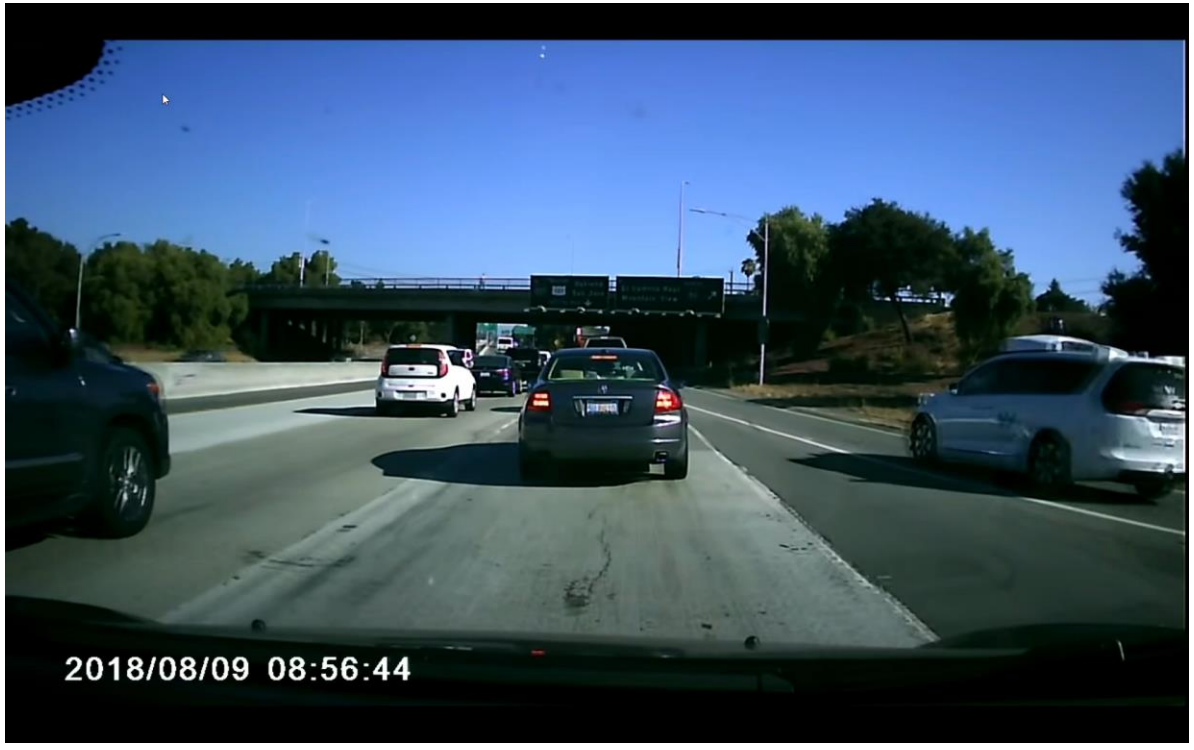
A. Zanardi (PhD Student @ IDSC Frazzoli)

Game-theoretical Motion Planning – Tutorial session

IEEE 2021 ICRA

Standard motion planning techniques lead to “passive” road users





- Resulting behaviours tend to be **overly cautious** and **passive**
- If predictions are ignored they become **hazardous**

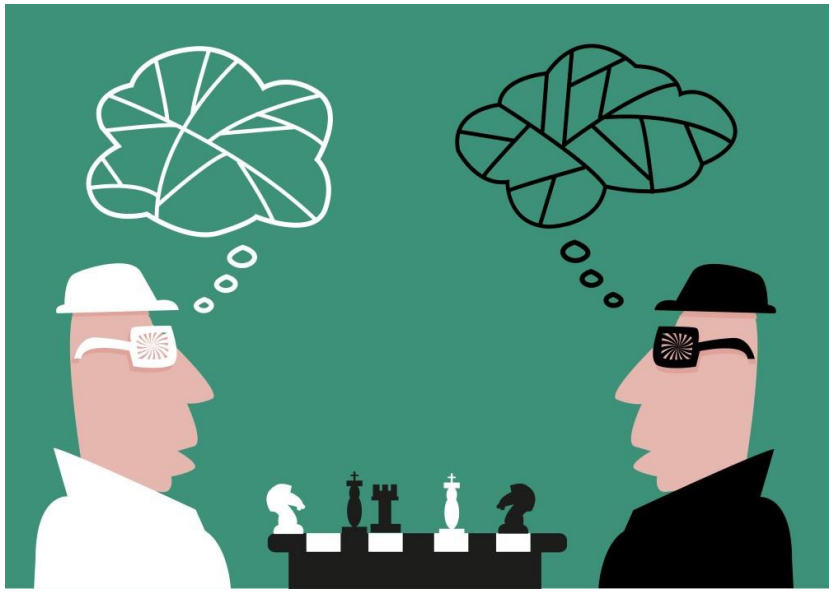
Explicitly taking into account others' decision making

Predicting what the others are about to do

Predicting what the others are about to do **conditioned on our actions**

Why game theory?

- Multiagent strategic decision making
- Most of the idealized assumptions can be relaxed
- Very little has been explored so far for **embodied agents**



Game theory applied to the driving task

- **Racing scenarios** [1,2,3], **highway merging** [4], **truck platooning** [7], common **intersections** [8],…

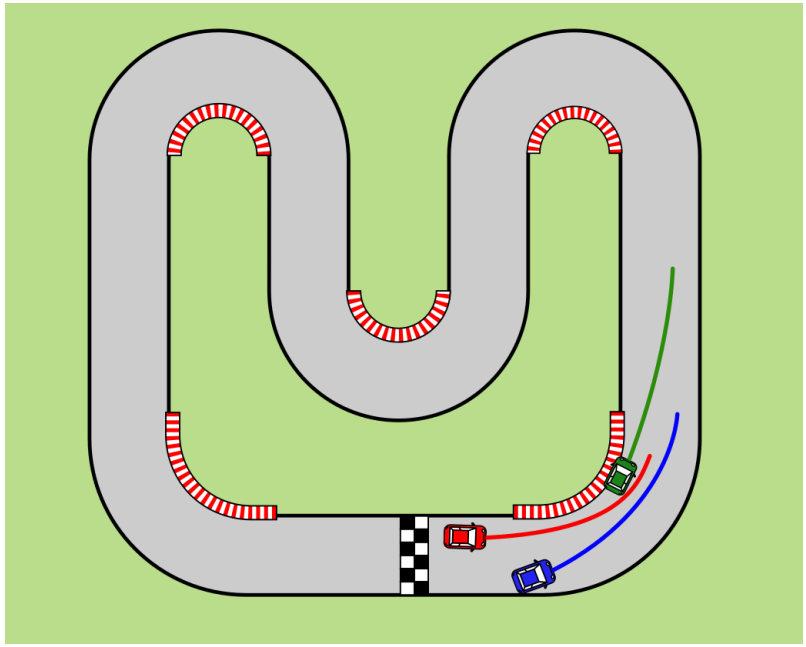
But the general game formulation is hard to crack:

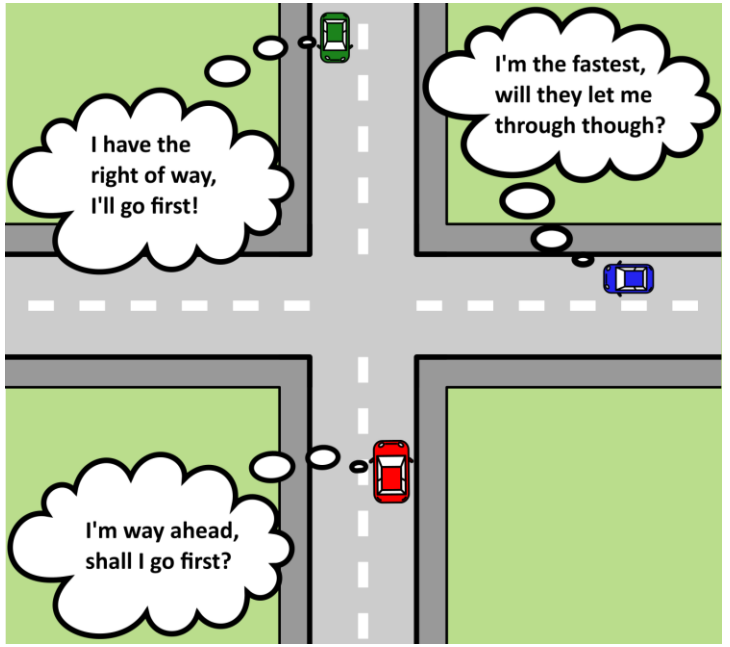
- Often using **Stackelberg solutions** [1,2,5], which assumes asymmetry of information among the players
- **Computationally quite demanding**, solutions are often limited to **2 players** [1,2,3,6], or to approximations [1,2,3,6]
- **Single scalar objective** [1,2,3,4,5,6,7,8] (Generalized Nash Equilibrium Problem [9])
- ...

[1] "*A Noncooperative Game Approach to Autonomous Racing*" – A. Liniger et al. (2019)
[2] "Game Theoretic Planning for Self-Driving Cars in Competitive Scenarios" – M. Wang et al. (2019)
[3] "A Real-Time Game Theoretic Planner for Autonomous Two-Player Drone Racing" – R. Spica et al. (2018)
[4] "ALGAMES: A Fast Solver for Constrained Dynamic Games" – S. Cleac'h et al. (2020)
[5] "Stackelberg Game Based Model of Highway Driving" - J. H. Yoo et al. (2012)
[6] "Hierarchical Game-Theoretic Planning for Autonomous Vehicles" J. Fisac et al. (2019)
[7] "Human-robot interaction for truck platooning using hierarchical dynamic games" - E. Stefansson et al. (2019)
[8] "Efficient Iterative Linear-Quadratic Approximations for Nonlinear Multi-Player General-Sum Differential Games" – D. Friedovich-Keil (2019)
[9] "Generalized Nash equilibrium problems" – F. Facchinei et al. (2007)
… many others

The game theory needed for urban driving

- **Not adversarial** (different from racing), **but non-cooperative** (agents are self-interested)





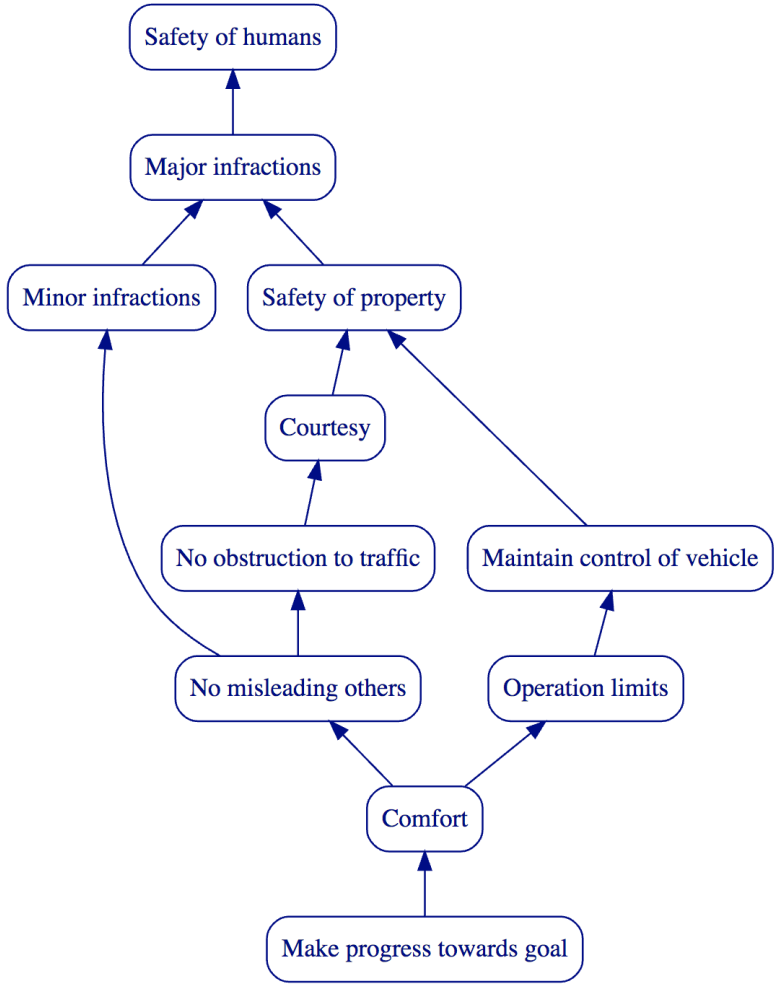
- Conceptually **zero-sum**,
  a winner (+1) and a loser (-1)
- **Adversarial scenario**,
  the optimization is a min/max problem

- **General-sum**,
  each player has personal interests
- The setup is **non-cooperative**
  but not adversarial

The game theory needed for urban driving

- **Multi-objective problem**, often conflicting

- **Decisions are** more **complex** than a scalar (e.g. Rulebooks):
  - Non trivial implications (e.g. liability)
  - Hard to do proper behaviour specifications with weighted sums



[*"Liability, Ethics, and Culture-Aware Behavior Specification using Rulebooks"* - A. Censi et al.]

UDGs with Lexicographic Preferences and Socially Efficient Nash Equilibria

In this work:

- We **describe** a precise structure of **Urban Driving Games (UDGs)**
- Few assumptions are needed to have a **potential structure**
- The potential structure of UDGs enjoys favorable properties:
  - **Socially efficient Nash Equilibria**
  - **Convergence guarantees for IBR-schemes**

Main diffence with related works:

Generalized Nash equilibrium problem (scalar cost function and constraints (e.g. collision))



- Agents express a **preference over** the possible **outcomes (non-scalar objectives)**
- The only constraints are the actuation limits

Autonomous agents interactions as differential games

Refresher on standard game theory

- The **state of the game** evolves according to a dynamic equation

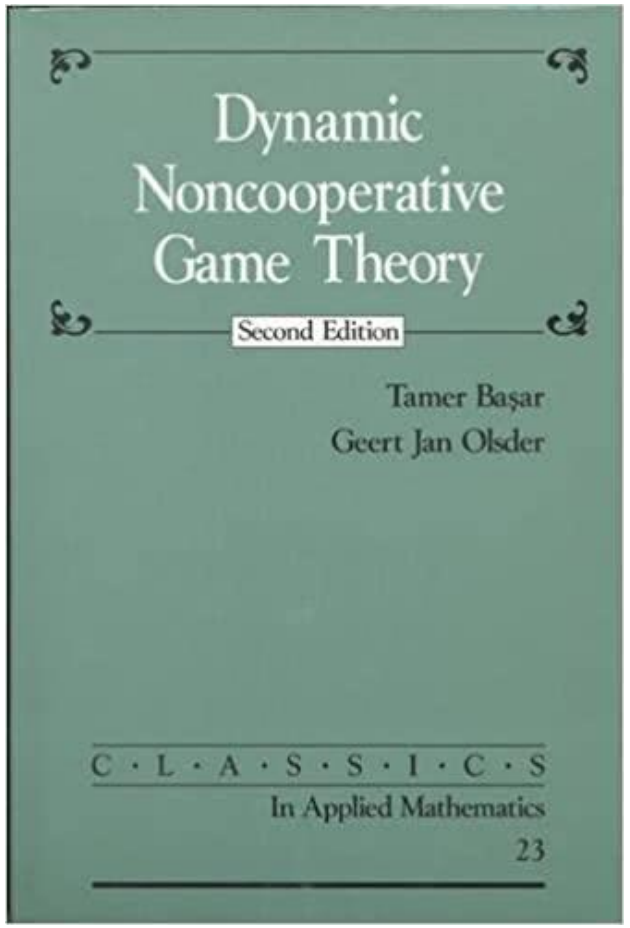
$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}_1(t), \ldots, \mathbf{u}_m(t), t)$$

- **Strategies** of players are function of the state (and time)

$$\mathbf{u}_i(t) = \gamma_i(t, \mathbf{x}(t)) \begin{cases} \gamma_i(\mathbf{x}(0)) & \text{open-loop} \\ \gamma_i(t, \mathbf{x}(t)) & \text{feedback strategies} \\ \vdots \end{cases}$$

- Costs of agents has an **incremental** part and a **terminal** component

$$J_i(\gamma_1, \ldots, \gamma_n) = \int_{t=0}^{T} g_i(t, \mathbf{x}(t), \mathbf{u}_1(t), \ldots, \mathbf{u}_n(t))dt + s_i(\mathbf{x}(T))$$



Autonomous agents interactions as differential games

Solutions

- Solution (**Nash Equilibria**)

$$J_i(\gamma_i^*, \gamma_{-i}^*) \leq J_i(\gamma_i, \gamma_{-i}^*) \qquad \forall \gamma_i \in \Gamma_i, \quad \forall i \in \mathcal{A}$$

"No player has an incentive to deviate"

- ***N*-coupled optimal control problems!** Very hard to solve in general:
  (that's why so far, a lot of approximations have been used)

  - **Minimum principle** with cross-terms in Hamiltonian costate dynamics, require partial derivatives wrt the equilibrium strategy.

  - **Dynamic programming** approach: solving *n*-coupled Hamilton-Jacobi-Bellman partial differential equations to find a set of value functionals (whose form needs to be guessed).

Games with structure are easier to solve though…

Common structure of Driving Games

- The **dynamics are decoupled**:
  Each agent influences only its own state.

$$x_i(k+1) = f_i(x_i(k), u_i(k))$$

- There is a **communal joint payoff** (*e.g. collision*) that depends **on the joint state and terminates the game.**

$$J_i^{\text{col}}(\gamma) = \sum_{j \in \{-i\}} J_{i,j}^{\text{col}}(x_i, x_j)$$

- **Personal payoff**:
  Incremental cost + personal terminating condition
  (*e.g. time, comfort, reached destination, race ends*)

$$J_i^{\text{per}}(\gamma) = \int_0^{T_i} g_i(x(t), u_i(t), t)dt + s_i(x_i(T_i))$$

- **Lexicographic preference:**
  The first concern is not to collide

$$J_i = \left\langle J_i^{\text{col}}(\gamma),\ J_i^{\text{per}}(\gamma) \right\rangle$$

Common structure of Driving Games

- There is a **communal joint payoff** (*e.g. collision*) that depends **on the joint state and terminates the game.**

$$J_i^{\text{col}}(\gamma) = \sum_{j \in \{-i\}} J_{i,j}^{\text{col}}(x_i, x_j)$$

Communal collision cost

"If it gets better for me,
it is better also for the others"

$$J_i^{\text{col}}(\gamma') < J_i^{\text{col}}(\gamma) \Rightarrow \sum_{j \in \{-i\}} J_{i,j}^{\text{col}}(\gamma') - J_{i,j}^{\text{col}}(\gamma') \leq 0$$

In practice, many common collision costs are allowed:

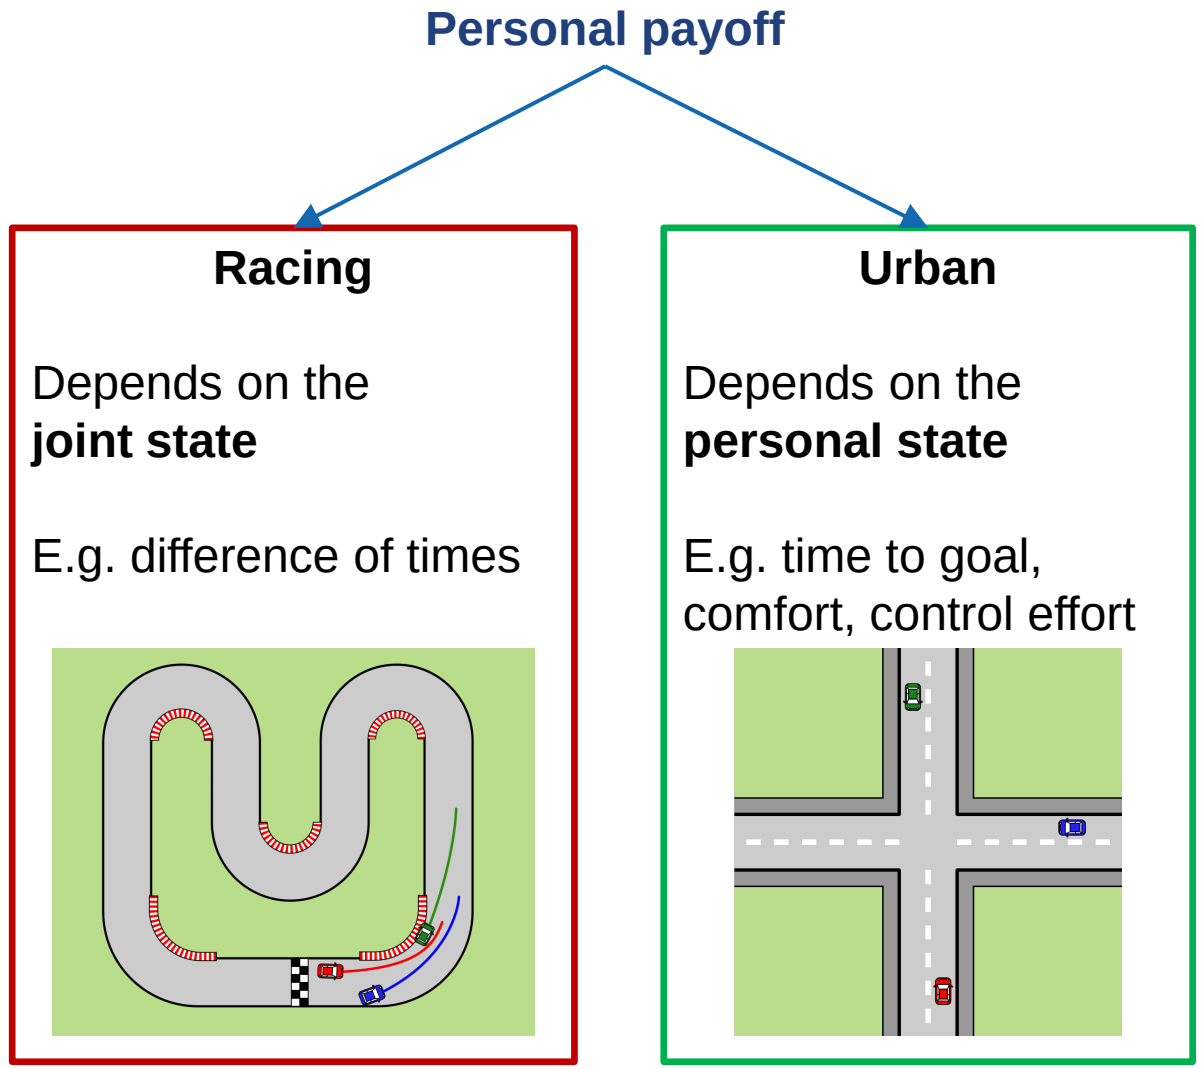
- Symmetric costs
- Same cost for everyone

E.g.: **kinetic energy transfer**, **indicator function**, **minimum safety distance**

Common structure of Driving Games

- **Personal payoff**:
  Incremental cost + personal terminating condition
  (*e.g. time, comfort + reached destination, race ends*)

$$J_i^{\text{per}}(\gamma) = \int_0^{T_i} g_i(x(t), u_i(t), t)dt + s_i(x_i(T_i))$$



Common structure of Driving Games

- **Lexicographic preference:**
  The first concern is not to collide

$$J_i = \left\langle J_i^{\text{col}}(\gamma),\ J_i^{\text{per}}(\gamma) \right\rangle$$

Lexicographic preference

$$J_i = \left\langle J_i^{\text{col}}(\gamma),\ J_i^{\text{per}}(\gamma) \right\rangle$$

1) **First**, do not collide
2) **Second**, minimize your personal cost

Note:

- Natural formalization of rational intents
- Better **expressiveness for behavior specification** (compared to weighted sums)
- Can be generalized to **arbitrary preference**

Communal Urban Driving Games

- The **dynamics are decoupled**:
  Each agent influences only its own state.

$$x_i(k+1) = f_i(x_i(k), u_i(k))$$

- There is a **communal joint payoff** (*e.g. collision*) that depends **on the joint state and terminates the game.**

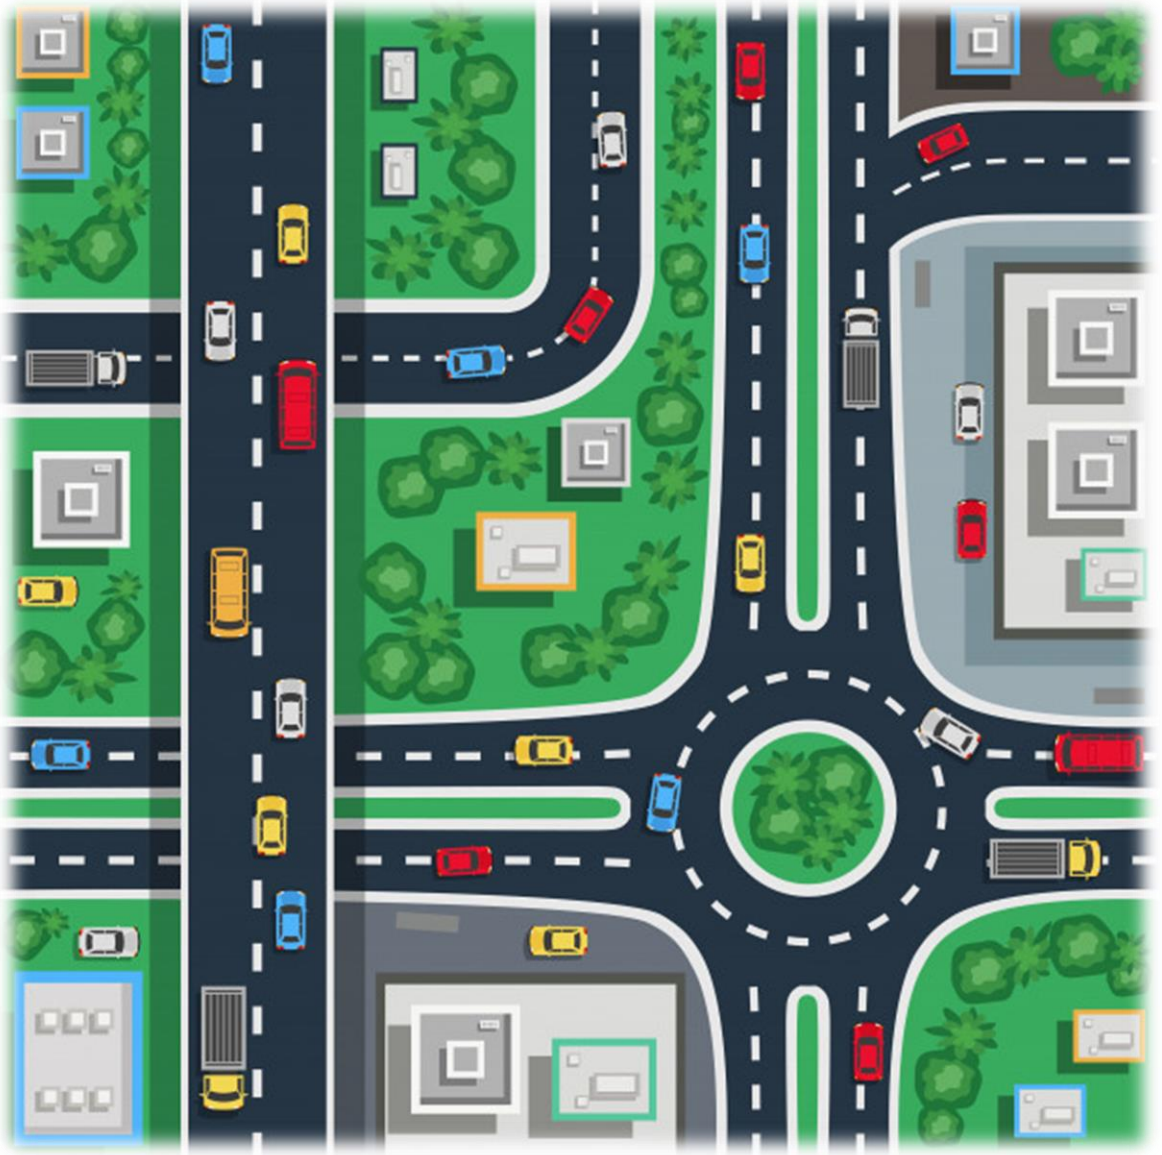
$$J_i^{\text{col}}(\gamma) = \sum_{j \in \{-i\}} J_{i,j}^{\text{col}}(x_i, x_j)$$

- **Personal payoff** that depends *only* on the personal state
  (*e.g. time, comfort + reached destination*)

$$J_i^{\text{per}}(\gamma) = \int_0^{T_i} g_i(x_i(t), u_i(t), t)dt + s_i(x_i(T_i))$$

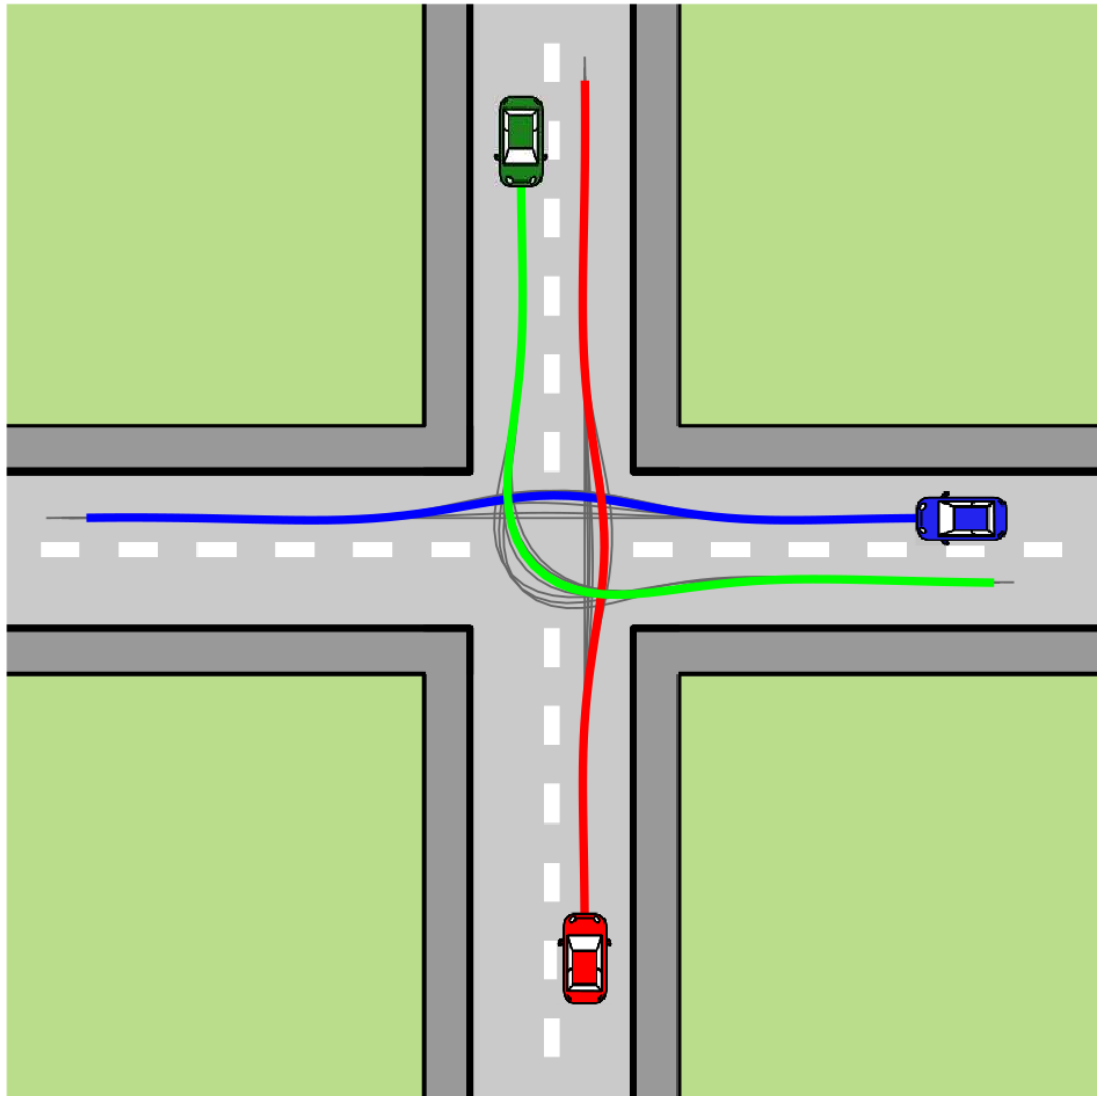
- **Lexicographic preference:**
  The first concern is not to collide

$$J_i = \left\langle J_i^{\text{col}}(\gamma),\ J_i^{\text{per}}(\gamma) \right\rangle$$



Social optima are Nash Equilibria in Urban Driving Games [1]

Theorem:
Under mild assumptions*, for an Urban Driving Game with communal collision cost it holds that the **minimum of the social cost is pure-strategy NE**.



* Compact strategy spaces and lower semi-continuous bounded functions.

Sketch of the proof (i)

- The statement follows from proving that a **Communal Urban Driving Game is a (lexicographic) potential** game

Def: A **potential game** admits a function $P : \times_i \Gamma_i \to S$ such that

$$J_i(\gamma_i', \gamma_{-i}) \prec J_i(\gamma_i'', \gamma_{-i}) \quad \Leftrightarrow \quad P(\gamma_i', \gamma_{-i}) \prec P(\gamma_i'', \gamma_{-i})$$

for every player i-th and every $\gamma_{-i} \in \Gamma_{-i}$, for every $\gamma_i', \gamma_i'' \in \Gamma_i$.

- A strategy profile minimizing the potential is a pure NE of the game (by def.)

We showed that **Urban Driving Games admit a potential function given by the social cost** (sum of individual costs)

(i) ✓

On IBR-schemes of CUDGs [1]

Proposition:
Let the assumptions of Thm. 1, then:

i. for an arbitrary small $\epsilon > 0$, **iterated-better response** (I-BR) **converges** to a Nash ϵ-equilibrium point **in a finite number of steps**.
ii. For discrete strategies, IBR-schemes converge to a NE.

Rational reaction set

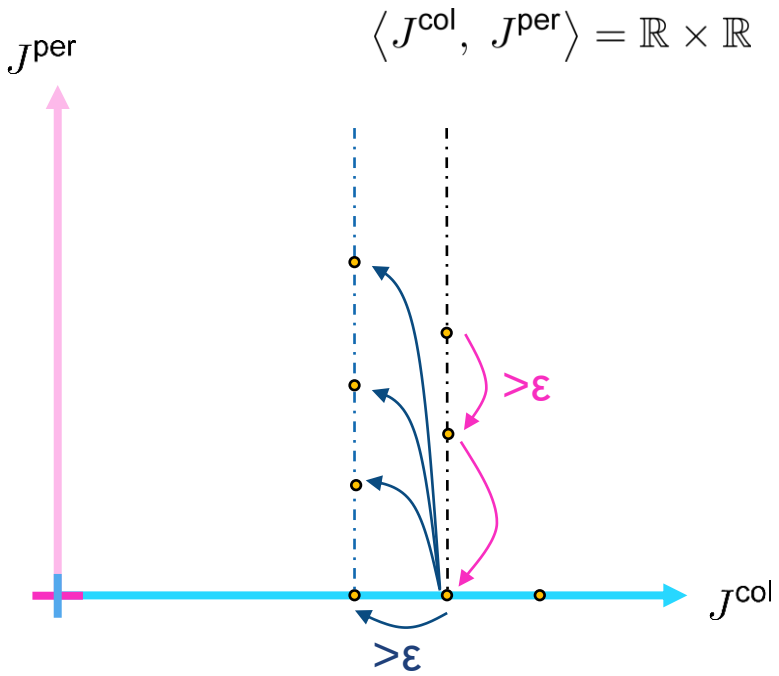
converged $\stackrel{?}{=} \gamma_i^k \in \mathcal{R}(\gamma_{-i}^k) \quad \forall i \in \mathcal{A}$;
$k = 0$;
$\gamma^0 = [\gamma_1^0, \ldots, \gamma_n^0]$;
while ¬converged **do**
  for $i \in \mathcal{A}$ **do**
    $\gamma_i^{k+1} \in \mathcal{R}_1(\gamma_{-i}^k)$
    $\gamma^{k+1} = [\gamma_1^k, \ldots, \gamma_i^{k+1}, \ldots, \gamma_n^k]$
    $k = k + 1$
  end
end

Note:

- Convergence of IBR to *"a"* NE
- For compact strategy spaces, one can only guarantee convergence to an ε-NE for an ε >0 arbitrary small. Yet at the limit (ε→0) one can have pathological cases for which best responses tend to a limit sequence whose limit point is not a NE.

Sketch of the proof (ii)

- Epsilon-**improvement on the lexicographic order:**



UDGs are potential games

⬇

If a player improves unilaterally by ε, also the potential improves by a finite amount

⬇

A myopic update of players strategies by “at least epsilon” (**iterated-better response**) **converges** in **finite number of steps** to an epsilon-Equilibrium

⬇

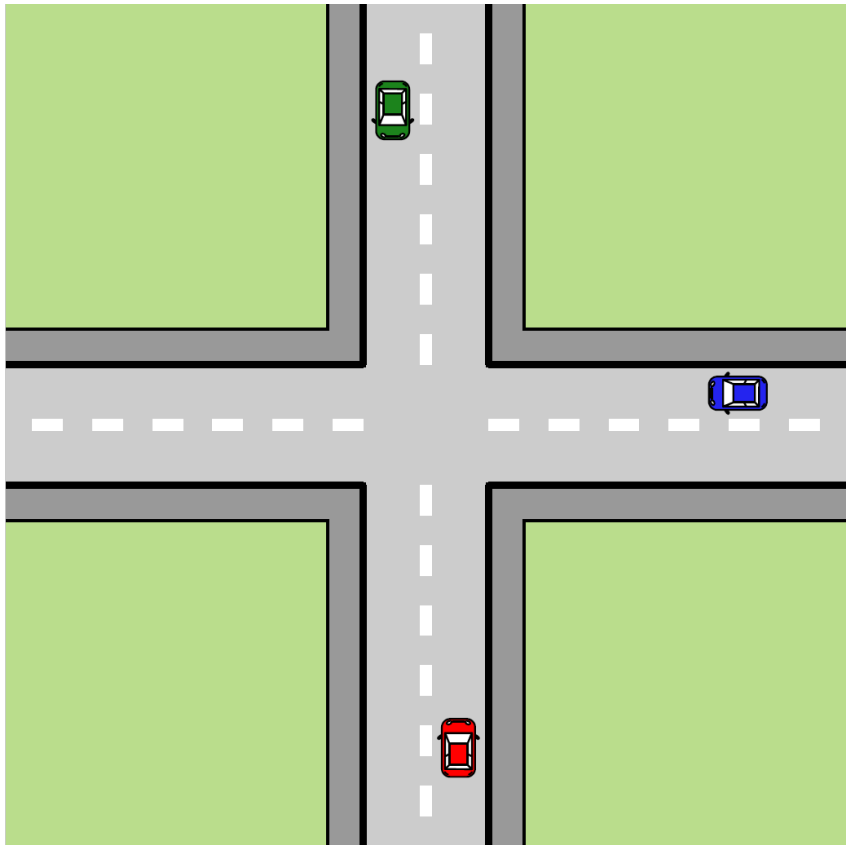
(ii) ✓

Practical relevance of the results

- A **pure NE is guaranteed to exist** (not true for general games)
- Pure NE are «**easy»** (as hard as a single non-linear optimization problem) **to compute**
  (*n* coupled player game can be solved as a multivariate **single** optimal control problem)
- Solving the OCP gives **socially efficient** equilibria
- **Best response iterations** (often instances of ε-better) **converge** to multiple NE
- Multiple NE found with IBR depend on initial conditions,
  the potential solution offers a ***rational*** **tool for equilibria refinement**

Study case: Trajectory Urban Games

Def: A **trajectory urban game** is an Urban Driving Game with **open loop information structure** (the agents commit to the whole trajectory).



- We compare NE trajectories obtained in 2 ways:

  1) **Iterated Best Response**
     (initialized with the individuals OCP)

  2) **Potential solution**

- Each agent's preference is a 3-levels lexicographic order

$$J_i = \left\langle J_i^{\text{col}}, J_i^{\text{per}(1)}, J_i^{\text{per}(2)} \right\rangle$$

Personal goal, comfort…

Traffic rules

Collision

Lexicographic optimization

- Easy for discrete strategies (compute metrics and sort trajectories)
- More challenging for continuous search spaces:

1) **Solve sequentially one optimization problem per lexicographic level**. The *k*-th optimization minimizes the *k*-th objective with the additional constraint of not worsening the previously computed costs.

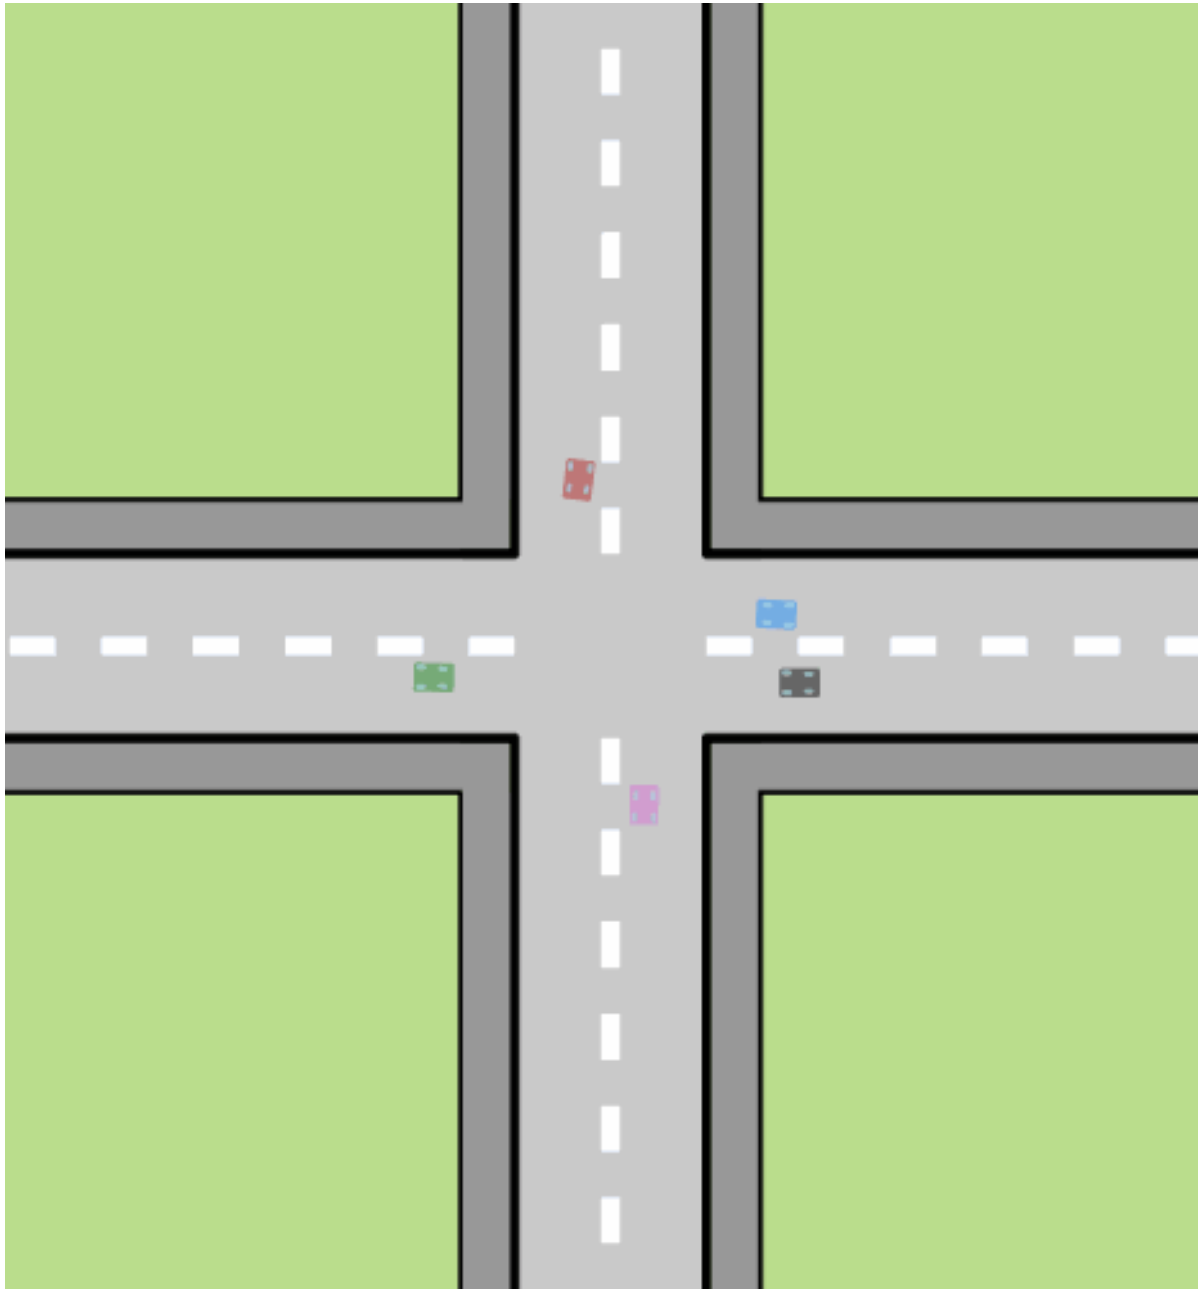
$$(J_i^{\text{col}})^* = \min_{\substack{\mathbf{u}_i \in U_i \\ \mathbf{x}_i(T_i) \in X_i^{\text{goal}}}} \quad J_i^{\text{col}}(\mathbf{x}_i, \mathbf{x}_{-i})$$
$$\text{s. t.} \quad \dot{\mathbf{x}}_i = f_i(\mathbf{x}_i, \mathbf{u}_i),\ \mathbf{x}_i^0$$

Followed by

$$\min_{\substack{\mathbf{u}_i \in U_i \\ \mathbf{x}_i(T_i) \in X_i^{\text{goal}}}} \quad J_i^{\text{per}}(\mathbf{x}_i, \mathbf{u}_i)$$
$$\text{s. t.} \quad J_i^{\text{col}}(\mathbf{x}_i, \mathbf{x}_{-i}) \leq (J_i^{\text{col}})^*$$
$$\dot{\mathbf{x}}_i = f_i(\mathbf{x}_i, \mathbf{u}_i),\ \mathbf{x}_i^0.$$

2) Solve the last optimization for different *"bottom"* values of higher rank objectives

Study case: Trajectory Urban Games

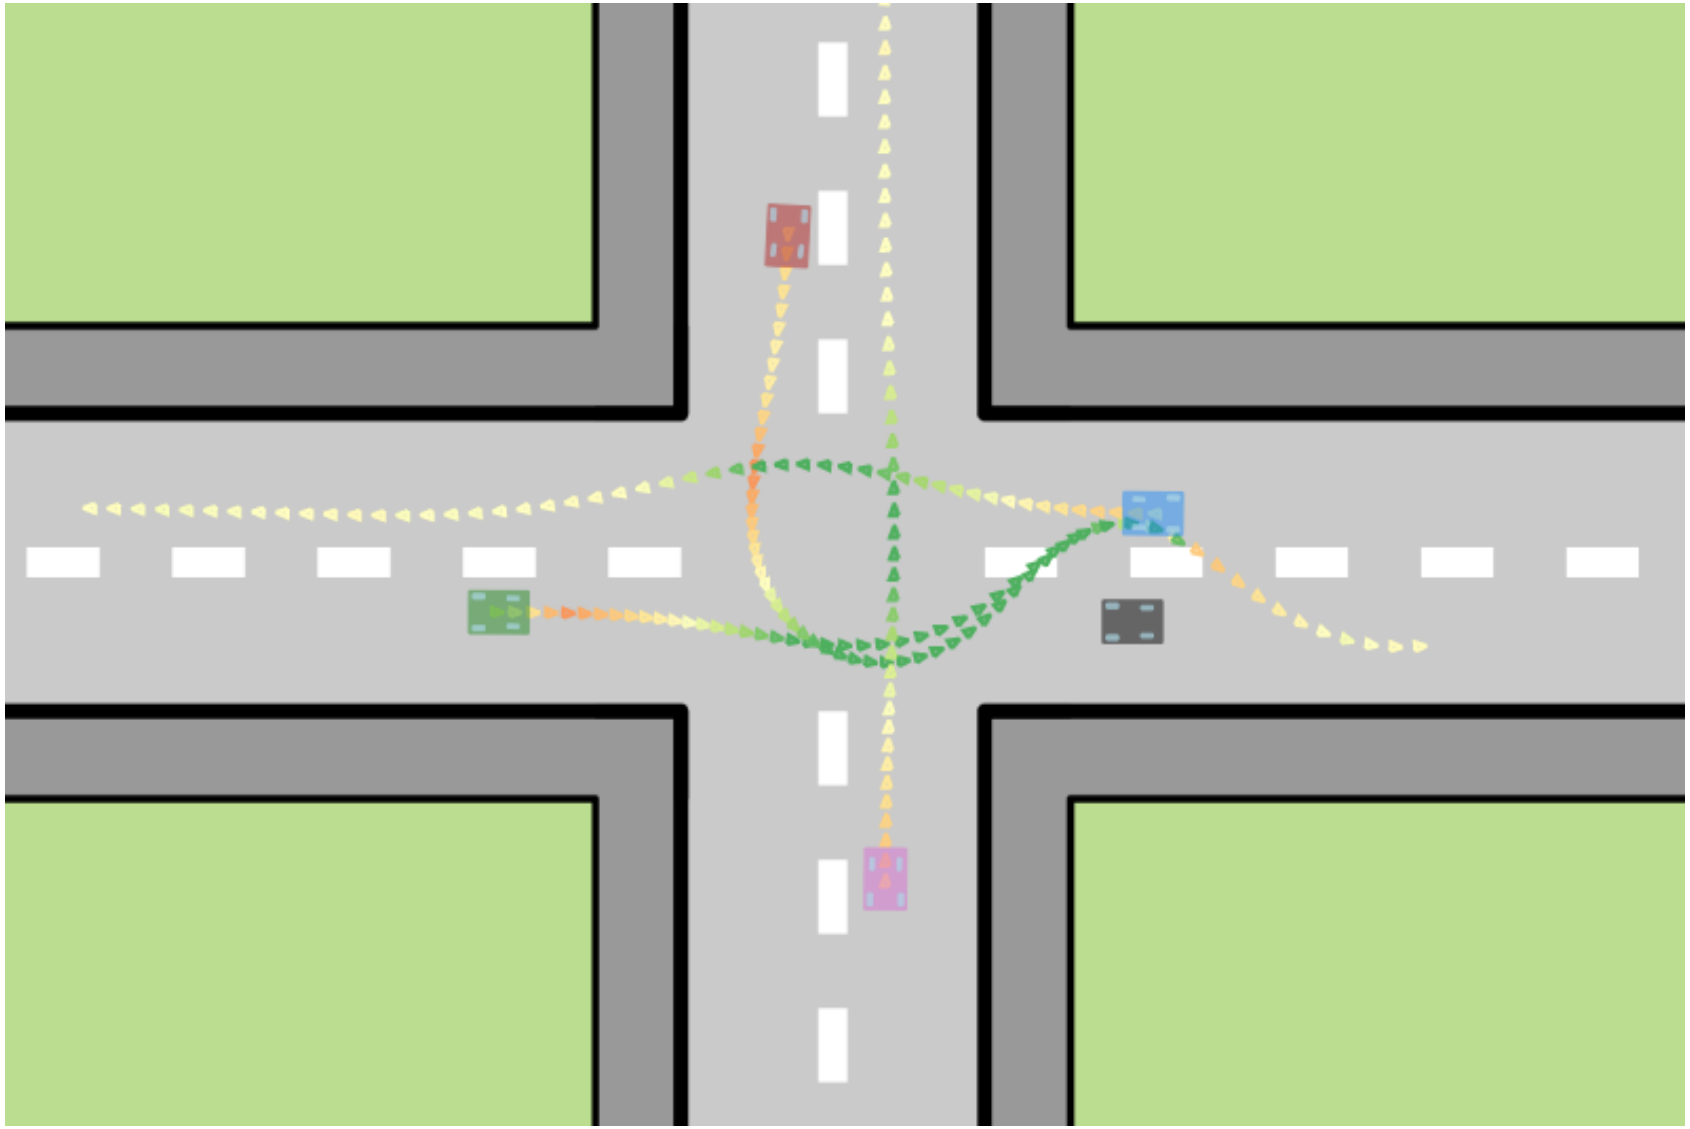


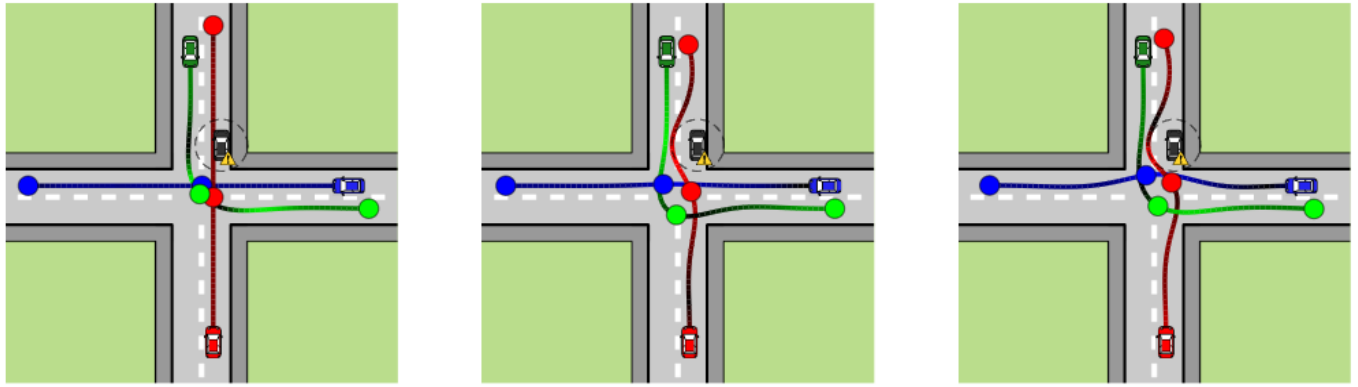
❖ 3-levels lexicographic preference:

$$J_i = \left\langle J_i^{\text{col}}, J_i^{\text{per}(1)}, J_i^{\text{per}(2)} \right\rangle$$

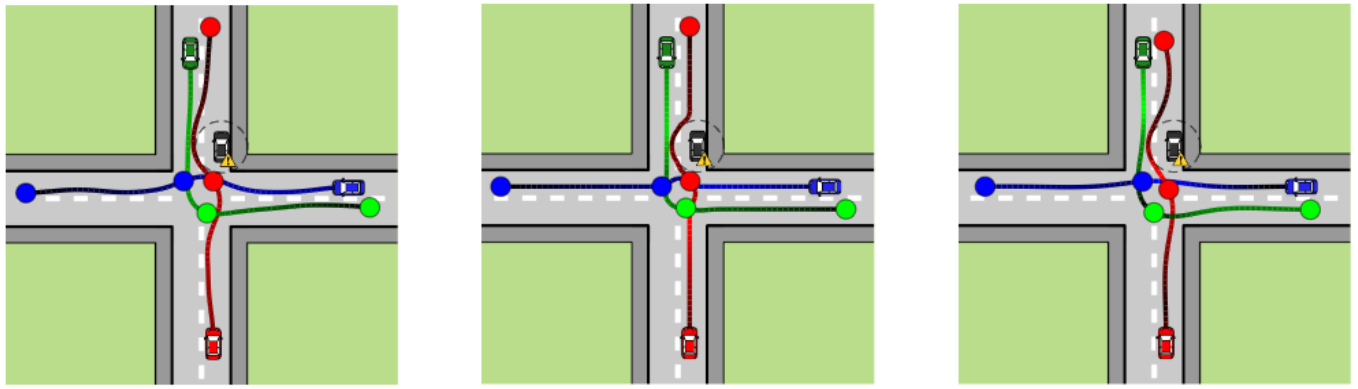
- $J_i^{\text{per}(2)}$: Personal goal, comfort…
- $J_i^{\text{per}(1)}$: Traffic rules
- J_i^{col}: Collision

Study case: Trajectory Urban Games

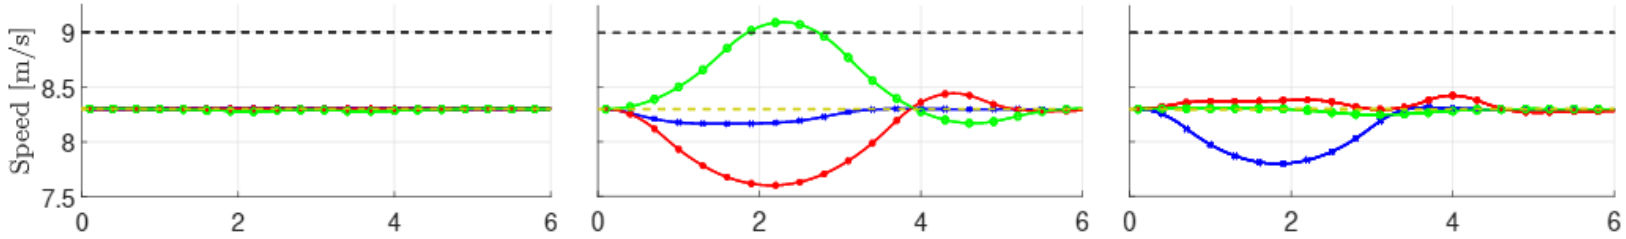




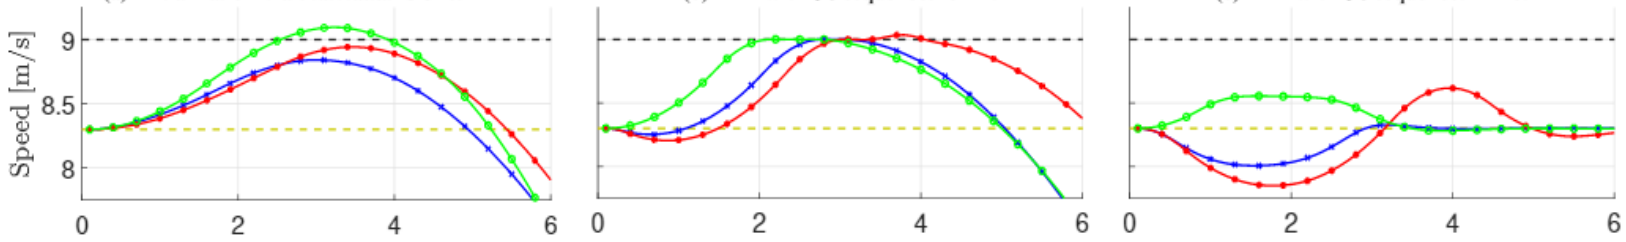
(a) Initialization via *individual* OCPs. (b) IBR with $\mathcal{A}$-sequence: GRB. (c) IBR with $\mathcal{A}$-sequence: RBG.



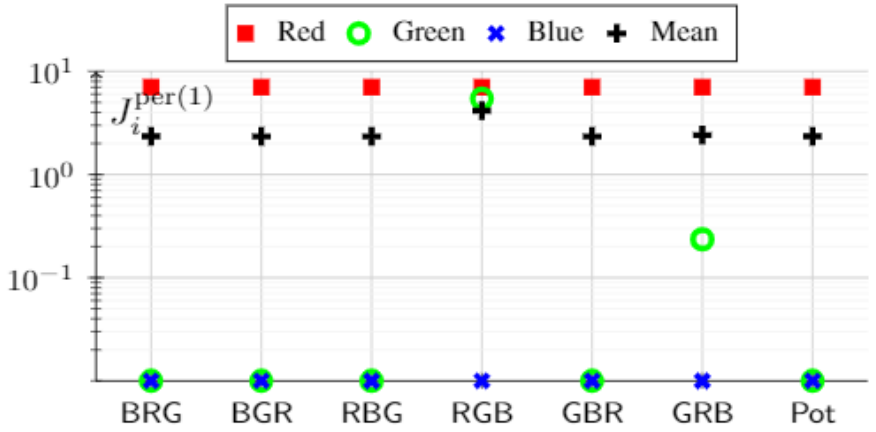
(d) 1^{st} lexicographic min of Pot (J^{col}). (e) 2^{nd} lexicographic min of Pot ($J^{\text{per}(1)}$). (f) 3^{rd} lexicographic min of Pot ($J^{\text{per}(2)}$).



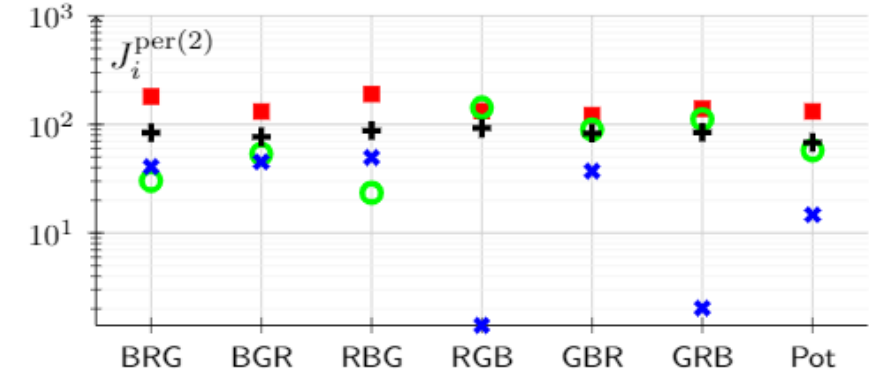
(a) Initialization via *individual* OCPs. (b) IBR with $\mathcal{A}$-sequence: GRB. (c) IBR with $\mathcal{A}$-sequence: RBG.



(d) 1^{st} lexicographic min of Pot (J^{col}). (e) 2^{nd} lexicographic min of Pot ($J^{\text{per}(1)}$). (f) 3^{rd} lexicographic min of Pot ($J^{\text{per}(2)}$).



(a) Costs associated with traffic rules infractions.



(b) Costs associated with personal costs of lower priority.

The potential solution is *the most rational*

Conclusion

- **Urban Driving Games (UDGs)** as general-sum games with a precise but comprehensive structure

- **Potential structure** of CUDGs:
  - **Theoretical guarantees** (Existence and efficiency of pure NE, convergence of IBR schemes)
  - **Computational implications** (n-OCPs → single-OCP)
  - **Rational tool for equilibria refinement**

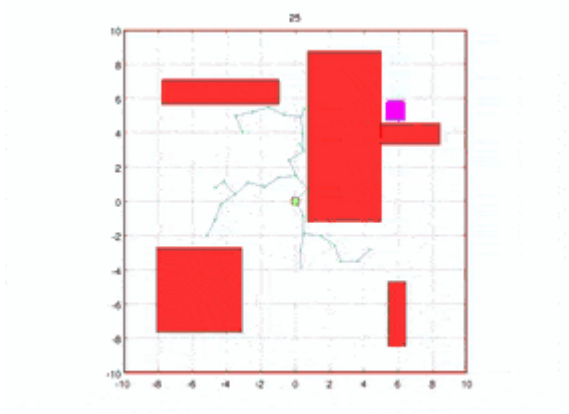
- **Preference relations on outcomes** to describe a richer decision making

Open Problems for UDGs

- Dealing with **many players**



- From kino-dynamic optimization to **high level planning** (complete algorithms,…)



- Including **bounded sensing and computation** (bounded rationality)



- Proper **inclusion of uncertainty** in the games (liability and intents)



Q&A

Contact: azanardi@ethz.ch