Hierarchical Decision-Making in Population Games

Yu-Wen Chen[1], Nuno C. Martins[2], Murat Arcak[1]

[1]Department of Electrical Engineering and Computer Sciences, UC Berkeley
[2]Department of Electrical and Computer Engineering, University of Maryland

December 11, 2025

paper/slides

Large-scale strategic multi-agent systems (MASs)


(a) transportation


(b) smart grids


(c) markets


(d) resource allocations


(e) Drone-as-a-Service

figures generated with Google Gemini

Large-scale strategic MASs: population games

- a large number of homogeneous players
- a strategy set $\mathcal{S} = \{1, \ldots, n\}$

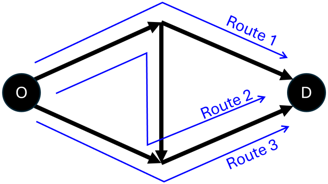

Large-scale strategic MASs: population games

- a large number of homogeneous players
- a strategy set $\mathcal{S} = \{1, \ldots, n\}$
- a social state $\boldsymbol{x} = (x_1, \ldots, x_n) \in \Delta^n := \{x_i \geq 0, \sum_i x_i = 1\}$

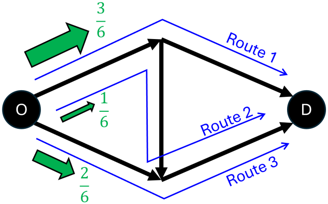

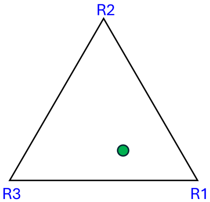

- a large number of homogeneous players
- a strategy set $\mathcal{S} = \{1, \ldots, n\}$
- a social state $\boldsymbol{x} = (x_1, \ldots, x_n) \in \Delta^n := \{x_i \geq 0, \sum_i x_i = 1\}$
- a payoff function $F : \Delta^n \to \mathbb{R}^n$ maps $\boldsymbol{x}$ to payoffs $[\ldots, F_i(\boldsymbol{x}), \ldots]^T$

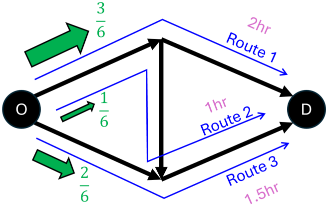

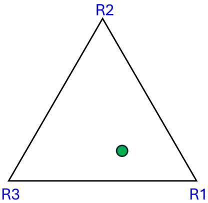

Large-scale strategic MASs: population games

- a large number of homogeneous players
- a strategy set $\mathcal{S} = \{1, \ldots, n\}$
- a social state $\boldsymbol{x} = (x_1, \ldots, x_n) \in \Delta^n := \{x_i \geq 0, \sum_i x_i = 1\}$
- a payoff function $F : \Delta^n \to \mathbb{R}^n$ maps $\boldsymbol{x}$ to payoffs $[\ldots, F_i(\boldsymbol{x}), \ldots]^T$
- an incentive-oriented revision protocol[1] $\mathcal{T} : \Delta^n \times \mathbb{R}^n \to \mathbb{R}_+^{n \times n}$, where $\mathcal{T}_{ij}$ is the switching rate from strategy i to j, inducing mean dynamics

$$\dot{x}_i = \sum_{j=1}^{n} x_j \mathcal{T}_{ji}(\boldsymbol{x}, F(\boldsymbol{x})) - x_i \mathcal{T}_{ij}(\boldsymbol{x}, F(\boldsymbol{x})), \quad i \in \mathcal{S},$$

which we call the evolutionary dynamics model (EDM)

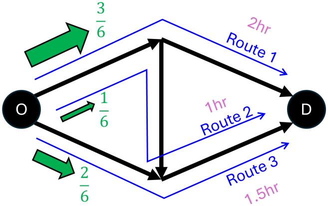

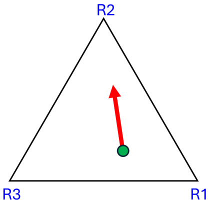

[1] Best Response, Pairwise Comparison, Imitation, Excess Payoff

Main studies in population games: Convergence

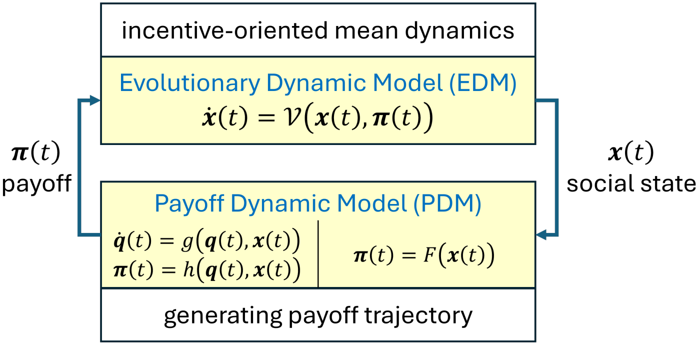

- with EDMs, what properties of PDMs lead to convergence?
  - not asking agents to follow some designed policies
  - not only about finding equilibrium, but about convergence of dynamics

Overview of our paper: hierarchical population games

- Motivation:
  Most population games assume that agents make direct decisions.

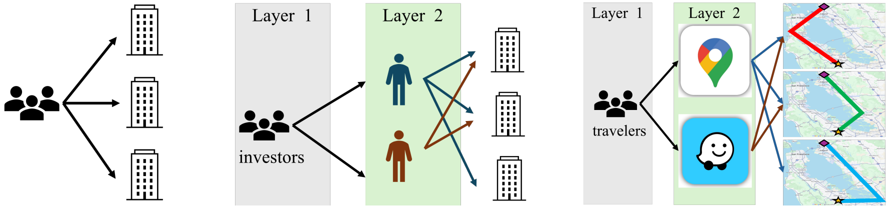

(a) direct framework (b) get info & make a decision (c) get info & make a decision

Overview of our paper: hierarchical population games

- Motivation:
  Most population games assume that agents make direct decisions.
- Contribution:
  - We propose a hierarchical decision-making framework for population games.
  - We prove convergence of system dynamics under some payoff structures.

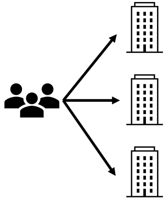

(a) direct framework

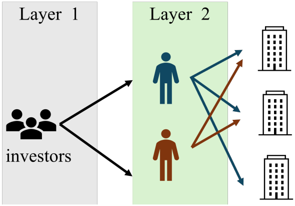

(b) get info & make a decision

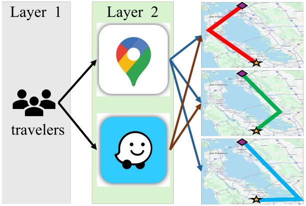

(c) get info & make a decision

Proposed basic hierarchical structure

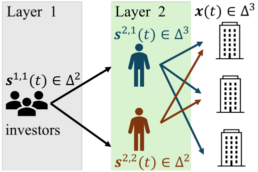

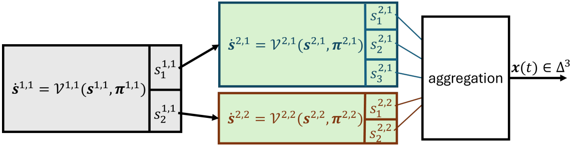

- each layer $i = 1, \ldots, L$ has n^i groups indexed by $j = 1, \ldots, n^i$
- $s^{i,j} \in \Delta^{d^{i,j}}$ is the distribution of (i,j)-group over its $d^{i,j}$ strategies

Proposed basic hierarchical structure

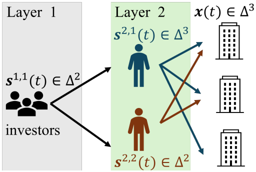

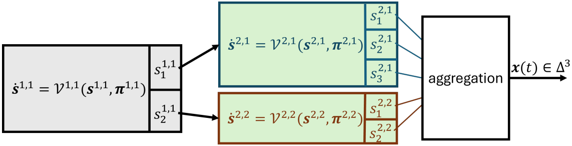

- each layer $i = 1, \ldots, L$ has n^i groups indexed by $j = 1, \ldots, n^i$
- $\boldsymbol{s}^{i,j} \in \Delta^{d^{i,j}}$ is the distribution of (i, j)-group over its $d^{i,j}$ strategies
- social state $\boldsymbol{x}(t) = \begin{bmatrix} s_1^{1,1}(t)s_1^{2,1}(t) + s_2^{1,1}(t)s_1^{2,2}(t) \\ s_1^{1,1}(t)s_2^{2,1}(t) + s_2^{1,1}(t)s_2^{2,2}(t) \\ s_1^{1,1}(t)s_3^{2,1}(t) \end{bmatrix}$: percentage to targets

Proposed basic hierarchical structure

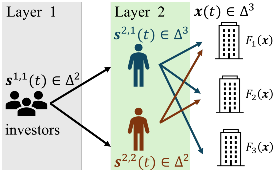

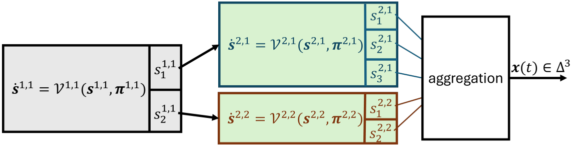

- each layer $i = 1, \ldots, L$ has n^i groups indexed by $j = 1, \ldots, n^i$
- $\boldsymbol{s}^{i,j} \in \Delta^{d^{i,j}}$ is the distribution of (i,j)-group over its $d^{i,j}$ strategies
- social state $\boldsymbol{x}(t) = \begin{bmatrix} s_1^{1,1}(t)s_1^{2,1}(t) + s_2^{1,1}(t)s_1^{2,2}(t) \\ s_1^{1,1}(t)s_2^{2,1}(t) + s_2^{1,1}(t)s_2^{2,2}(t) \\ s_1^{1,1}(t)s_3^{2,1}(t) \end{bmatrix}$: percentage to targets
- payoff $\boldsymbol{\pi}^{i,j}(t)$ is (backwardly) determined by the average performance

$\boldsymbol{\pi}^{2,1}(t) = F(\boldsymbol{x}(t))$, $\quad F(\cdot)$: social payoff function

$\boldsymbol{\pi}^{2,2}(t) = [F_1(\boldsymbol{x}(t)), F_2(\boldsymbol{x}(t))]^T$, $\quad F(\boldsymbol{x}) = [F_1(\boldsymbol{x}), F_2(\boldsymbol{x}), F_3(\boldsymbol{x})]^T$

$$\boldsymbol{\pi}^{1,1}(t) = \left[\sum_{l=1}^{3} s_l^{2,1}(t)F_l(\boldsymbol{x}(t)), \sum_{l=1}^{2} s_l^{2,2}(t)F_l(\boldsymbol{x}(t))\right]^T$$

- We consider general cases: $\boldsymbol{s}^{i,j} \in K^{i,j} \subseteq \Delta^{d^{i,j}}$, e.g.,
  - Manager 1 group is restricted to the distribution $\boldsymbol{s}_*^{2,1} = [0.2, 0.2, 0.6]^T$ over the three investment targets but allows variations up to $\varepsilon = 0.1$,

$$K^{2,1} = \left\{ \boldsymbol{x} \in \Delta^3 : \left\| \boldsymbol{x} - \boldsymbol{s}_*^{2,1} \right\|_2 \leq \varepsilon \right\}.$$

  - Manager 2 group needs to invest at least $1/3$ in both investment targets,

$$K^{2,2} = \left\{ \boldsymbol{x} \in \Delta^2 : x_1 \geq \frac{1}{3}, x_2 \geq \frac{1}{3} \right\}.$$

$\Rightarrow$ social state $\boldsymbol{x}(t)$ is confined within the admissible set K.

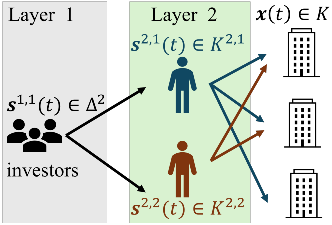

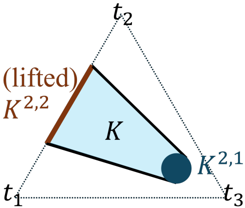

Research focus on hierarchical population games

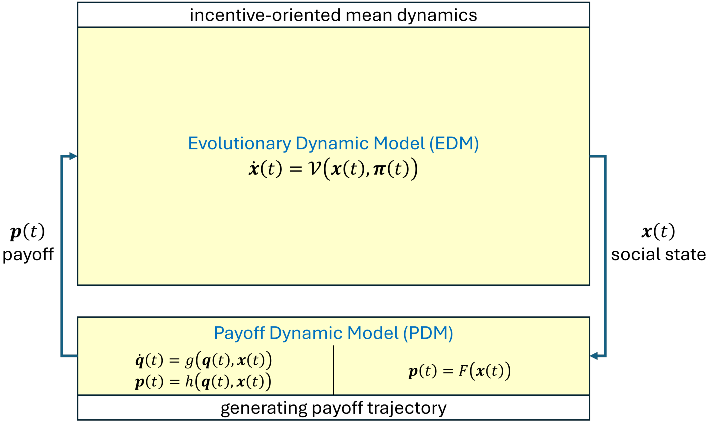

Research focus on hierarchical population games

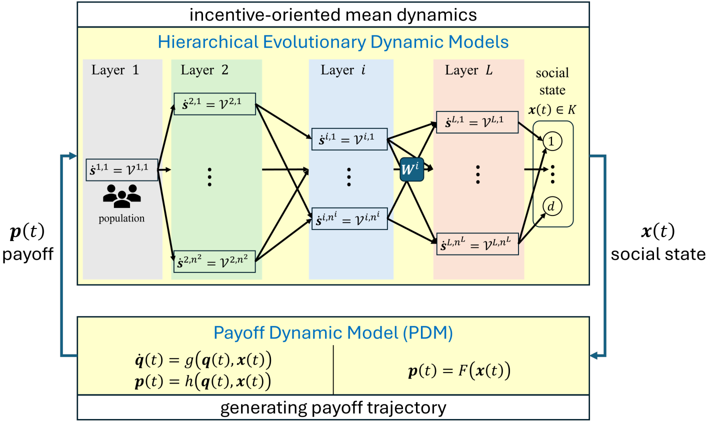

- properties on PDM s.t. $\boldsymbol{x}(t)$ evolves within K and converges
  - in classical case: agents react directly to the payoff function F
    $\Rightarrow$ rest points correspond to the set of Nash equilibria NE(F)

Research focus on hierarchical population games

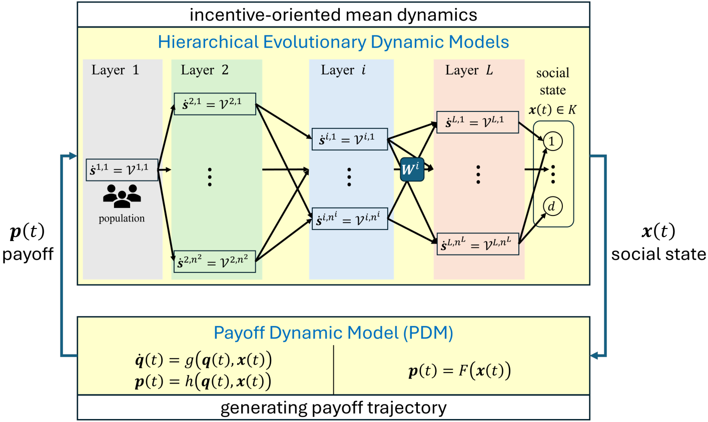

- properties on PDM s.t. $\boldsymbol{x}(t)$ evolves within K and converges
  - in classical case: agents react directly to the payoff function F
    $\Rightarrow$ rest points correspond to the set of Nash equilibria $\mathrm{NE}(F)$
  - in hierarchical case: rest points meaningful? related to $\mathrm{NE}_K(F)$?

Existence of $\mathrm{NE}_K(F)$

Theorem (Convexity of K)

If $K^{i,j}$s are all convex, then K is convex.

$\Rightarrow$ By Kakutani's Fixed-Point Theorem:
The convexity of K guarantees the existence of $\mathrm{NE}_K(F)$.

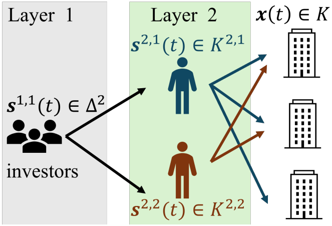

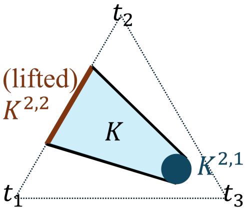

Rest points are (refined) Nash equilibria

Characterization at rest points (informal)

When each $\boldsymbol{s}^{i,j}$ attains a Nash equilibrium in its population game (w.r.t. $\boldsymbol{\pi}^{i,j}$), the corresponding social state $\boldsymbol{x}$ satisfies $\boldsymbol{x} \in \mathsf{NE}_K(F)$.

- all groups ($\boldsymbol{s}^{i,j}$) contribute to the social state ($\boldsymbol{x}$)
- all groups are playing population games w/ different payoff functions

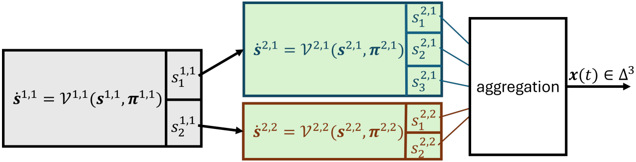

Rest points are (refined) Nash equilibria

Characterization at rest points (informal)

When each $\boldsymbol{s}^{i,j}$ attains a Nash equilibrium in its population game (w.r.t. $\boldsymbol{\pi}^{i,j}$), the corresponding social state $\boldsymbol{x}$ satisfies $\boldsymbol{x} \in \mathsf{NE}_K(F)$.

- all groups ($\boldsymbol{s}^{i,j}$) contribute to the social state ($\boldsymbol{x}$)
- all groups are playing population games w/ different payoff functions

$\Rightarrow$ the proposed average payoff determination for $\boldsymbol{\pi}^{i,j}$ renders that the decision groups, as players, participate in a (weighted) potential game

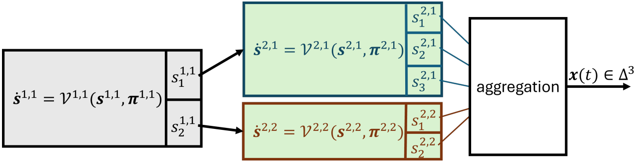

Convergence results for hierarchical population games

Definition (Nash Stationarity (NS) & Positive Correlation (PC))

An EDM is
(1) NS w.r.t K if $\mathcal{V}(\boldsymbol{s}, \boldsymbol{\pi}) = \mathbf{0} \iff \boldsymbol{s} \in \beta_K(\boldsymbol{\pi}) = \arg\max_{\boldsymbol{y} \in K} \boldsymbol{y}^T \boldsymbol{\pi}$.
(2) PC if $\mathcal{V}(\boldsymbol{s}, \boldsymbol{\pi}) \neq \mathbf{0} \implies \boldsymbol{\pi}^T \mathcal{V}(\boldsymbol{s}, \boldsymbol{\pi}) > 0$.

Theorem (Convergence for potential game F)

Let F be a potential game, and some regularities hold (Appendix). Suppose each $\mathcal{V}^{i,j}$ is NS w.r.t. $K^{i,j}$ and PC. Then, $\boldsymbol{x}(t)$ asymptotically approaches $\text{NE}_K(F)$.

Theorem (Convergence for counterclockwise dissipativity PDM)

Let the payoff $\boldsymbol{p}(t)$ be given by a counterclockwise dissipativity PDM. Suppose that each $\mathcal{V}^{i,j}$ is NS w.r.t. $K^{i,j}$ and PC. Then, $\boldsymbol{x}(t) \to \beta_K(\boldsymbol{p}(t))$.

Conclusions

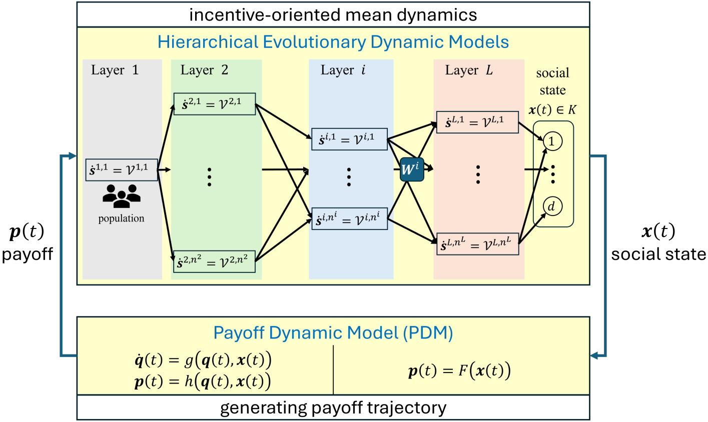

paper/slides

website

- Proposed a hierarchical decision-making structure for population games.
- Proved convergence under some payoff structures.

Appendix

Assumptions

1. $K^{i,j}, i \in [L], j \in [n^i]$, are compactly convex.
2. F is continuously differentiable.
3. Each $\mathcal{V}^{i,j}$ ensures the forward invariant of $K^{i,j}$.
4. $\boldsymbol{\pi}^T \mathcal{V}^{i,j}(\boldsymbol{s}, \boldsymbol{\pi})$, $i \in [L], j \in [n^i]$, are Lipschitz continuous w.r.t. $\boldsymbol{s}$ and $\boldsymbol{\pi}$.

$\Rightarrow$ When $L = 1$ and $K^{1,1} = \Delta^{d^{1,1}}$, the assumptions reduce to the ones in classical population games.

$\Rightarrow$ Assumption 3 means that each (i, j)-group makes decisions consistently within its strategy distribution of interest $K^{i,j}$.

As a design tool: partial information revealing

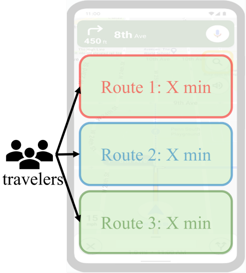

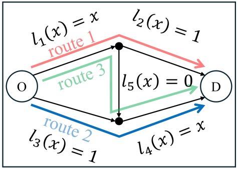

- If directly provide all information, then travelers face the Braess paradox.

As a design tool: partial information revealing

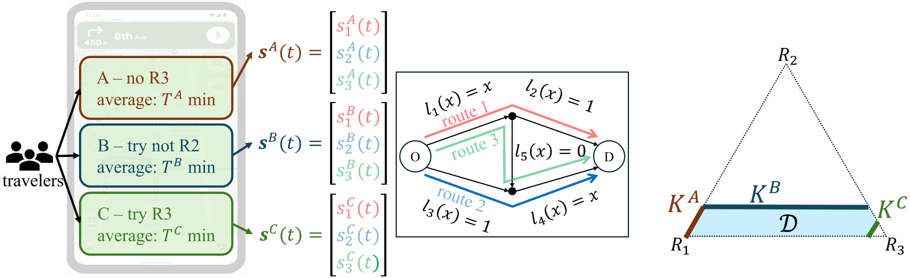

- If directly provide all information, then travelers face the Braess paradox.
- Instead, provide (verifiable) partial information, like:
  - each option plays a population game for users based on its known information
  - joining through option X undergoes average time T^X, for $X = A, B, C$

As a design tool: partial information revealing

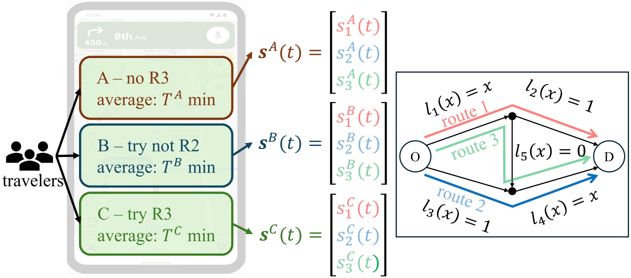

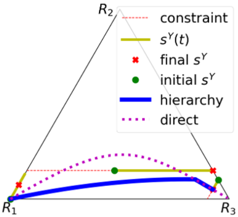

- If directly provide all information, then travelers face the Braess paradox.
- Instead, provide (verifiable) partial information, like:
  - each option plays a population game for users based on its known information
  - joining through option X undergoes average time T^X, for $X = A, B, C$

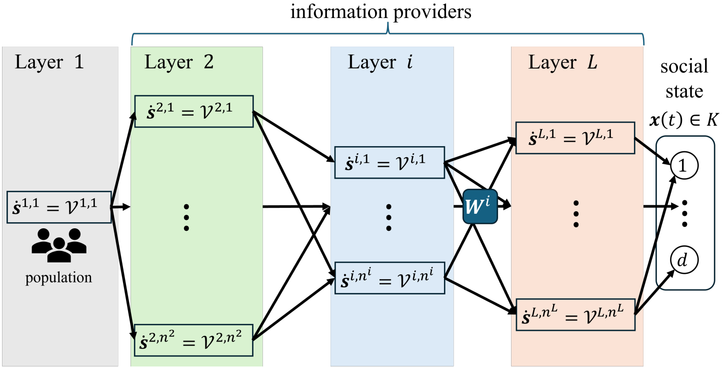

- agent level: revision protocol
- information providers: cooperation / competition