Min-LQG Games and Collective Discrete Choice Problems

Presented by Roland P. Malhamé
Joint work with Rabih Salhab and Jérôme Le Ny
GERAD and École Polytechnique de Montréal

MEAN FIELD GAMES AND RELATED TOPICS - 4
Chiostro di S. Pietro in Vincoli, Rome

June 14-16, 2017

Rabih Salhab

Jérôme Le Ny

Overview

Example 1 - "Tupperware Effect"

[Harvard Gazette (2012) - "Peer pressure in elections"]

- **Event:** Fundraising coffee session for a 2009 Democratic U.S. House candidate, Julie Hamos.
- **Previous elections cyles:** The group supported another candidate.
- **2009:** After Julie's speech, the guests were invited to write checks, and many backed Julie.
- Julie *lost* and many guests *regretted* backing her.

What made a guest support Julie although initially not in favor of her candidacy?

Fear of being different from the group: Conformity pressure, Social effect, Tupperware Effect, Peer effect.

Example 2 - Social effect and the diffusion of home computers

A study by Goolsbe and Klenow [Goolsbe and Klenow '02].

- Samples of 208 U.S. cities.
- $x-$axis $=$ Fraction of households owning a computer at the beginning of 1997.
- $y-$axis $=$ Fraction of households not owning a computer at the beginning of 1997 that bought during 1997.

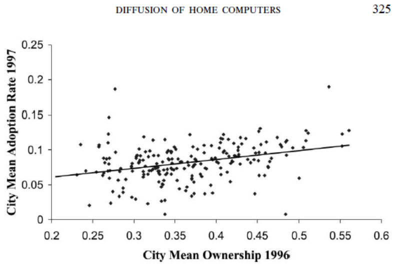

FIGURE 2.—City mean adoption rate

- **What are the collective choice problems ?**
  The collective choice problems are concerned with situations and decision making when a large number of agents make a socially influenced choice amongst different alternatives.

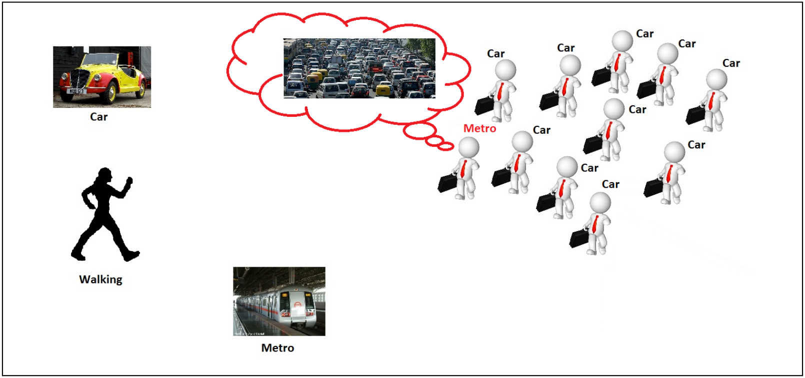

Application 1 - Robotics

- Collection of robots exploring an unknown terrain.
- Multiple potential sites of interest to visit (**Discrete choices**).
- The robots must stay as much as possible grouped to carry out some collective tasks (**Social effect**).

Application 2 - Elections

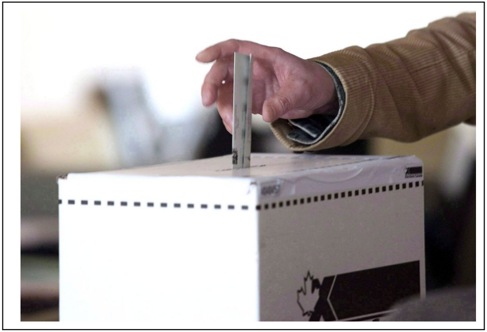

- A group of voters are choosing among a set of candidates (**Discrete choices**).
- Along the path to choose a candidate, changing one's opinion requires an effort but deviation from the majority's opinion involves a discomfort (**Social effect**).

Problem Statement

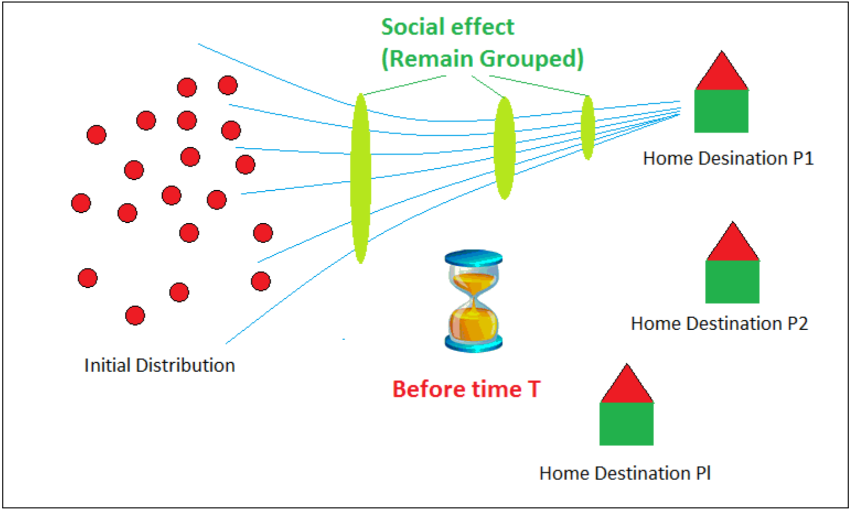

Questions We Addressed

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?

Questions We Addressed

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?

- Can one anticipate the macroscopic behavior of the population, i.e. the distribution of the agents between the alternatives?

Questions We Addressed

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
- Can one anticipate the macroscopic behavior of the population, i.e. the distribution of the agents between the alternatives?
- If yes, are there multiple potential macroscopic behaviors?

Questions We Addressed

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
- Can one anticipate the macroscopic behavior of the population, i.e. the distribution of the agents between the alternatives?
- If yes, are there multiple potential macroscopic behaviors?
- What is the necessary information to compute the individual strategies and anticipate the macroscopic behaviors?

Questions We Addressed

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
- Can one anticipate the macroscopic behavior of the population, i.e. the distribution of the agents between the alternatives?
- If yes, are there multiple potential macroscopic behaviors?
- What is the necessary information to compute the individual strategies and anticipate the macroscopic behaviors?
- How does the strength of the social effect influence the macroscopic behavior?

Min-LQG Model

N agents, with dynamics and individual costs:

$$dx_i = Ax_i dt + Bu_i dt + \sigma dw_i \qquad \forall i \in \{1 \dots, N\}, \tag{1}$$

$$J_i(u_i, \bar{x}, x_i^0) = \mathbb{E}\Bigg(\int_0^T \left\{ \frac{q}{2} \overbrace{\|x_i - \bar{x}\|^2}^{\text{Social effect}} + \frac{r}{2} \overbrace{\|u_i\|^2}^{\text{Effort}} \right\} \mathrm{d}t + \frac{M}{2} \min_{j=1,\dots,l} \Big(\overbrace{\|x_i(T) - p_j\|^2}^{\text{Choice}} \Big) \Big| x_i^0 \Bigg), \tag{2}$$

- $x_i \in \mathbb{R}^n$ the state of agent i, $u_i \in \mathbb{R}^m$ its control input.
- $p_j \in \mathbb{R}^n$ for $j = 1, \dots, l$ home destinations.
- $\bar{x} = 1/N \sum_{j=1}^N x_j$ the average of the population.
- Independent initial conditions x_i^0 and Wiener processes w_i, $i = 1, \dots, N$.

Methodology - Mean Field Games

For a continuum of agents

Assume $\bar{x}$ given and equal to $\hat{x}$

Methodology - Mean Field Games

For a continuum of agents

Assume $\bar{x}$ given and equal to $\hat{x}$ →

Methodology - Mean Field Games

For a continuum of agents

Assume $\bar{x}$ given and equal to $\hat{x}$ → Minimize $J_i(u_i, \hat{x})$ (**Tracking Problem**)

For a continuum of agents

Assume $\bar{x}$ given and equal to $\hat{x}$ → Minimize $J_i(u_i, \hat{x})$ (**Tracking Problem**) →

Methodology - Mean Field Games

For a continuum of agents

Assume $\bar{x}$ given and equal to $\hat{x}$ → Minimize $J_i(u_i, \hat{x})$ (**Tracking Problem**) → Decentralized strategies $u_i(x_i, \hat{x})$

Methodology - Mean Field Games

For a continuum of agents

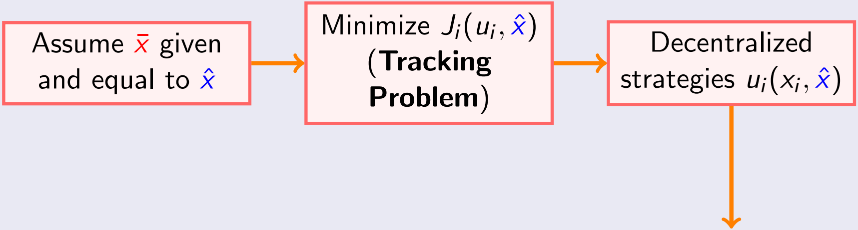

Methodology - Mean Field Games

For a continuum of agents

Assume $\bar{x}$ given and equal to $\hat{x}$

→ Minimize $J_i(u_i, \hat{x})$ (**Tracking Problem**)

→ Decentralized strategies $u_i(x_i, \hat{x})$

→ $dx_i = Ax_i dt + Bu_i dt$ (**State Aggregation**)

Methodology - Mean Field Games

For a continuum of agents

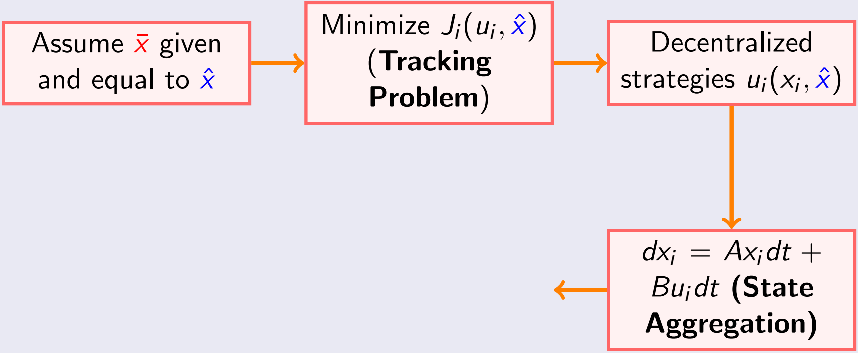

For a continuum of agents

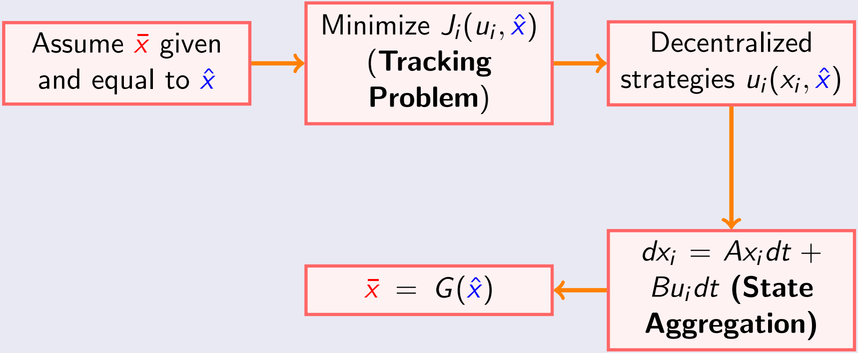

Methodology - Mean Field Games

For a continuum of agents

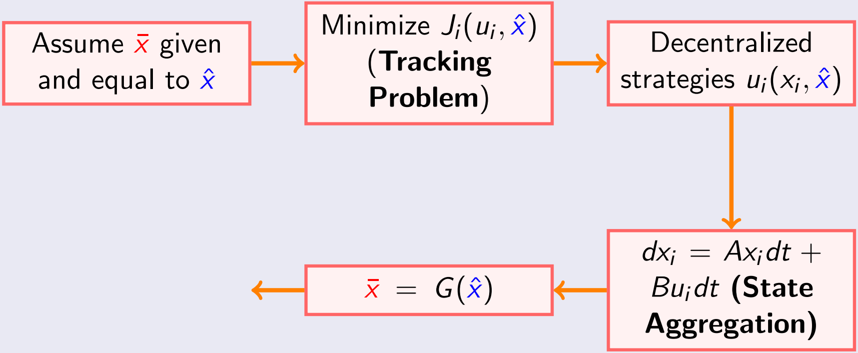

Methodology - Mean Field Games

For a continuum of agents

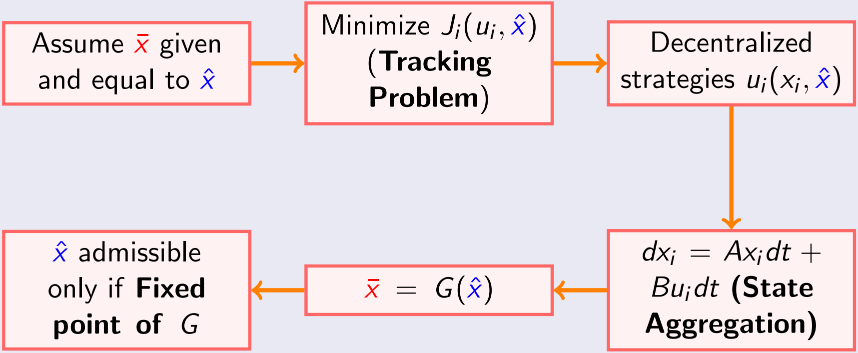

Methodology - Mean Field Games

For a continuum of agents

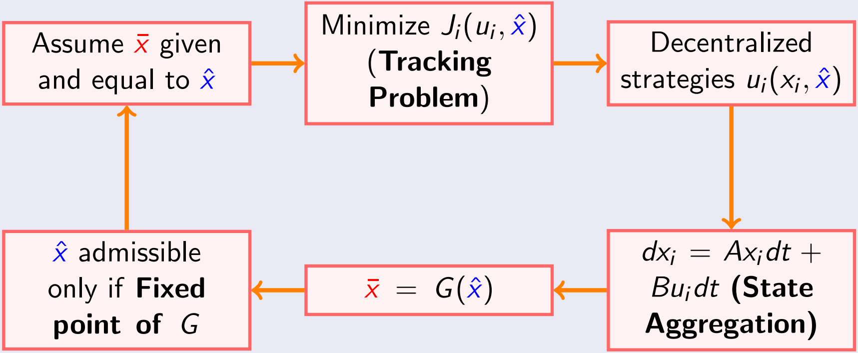

Methodology - Mean Field Games

Practical situation (finite number of agents)

The decentralized strategies $u_i(x_i, \hat{x})$ computed for a continuum of agents (hopefully) constitute an $\epsilon-$Nash equilibrium, where ϵ is close to zero for N large enough.

Definition

An $\epsilon-$Nash equilibrium is such that each agent can profit at most ϵ by a unilateral deviant behavior.

Main Contributions

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?

Main Contributions

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
  - **Without noise**: an agent picks a destination prior start moving based on its initial condition and the initial distribution of the population.

Main Contributions

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
  - **Without noise**: an agent picks a destination prior start moving based on its initial condition and the initial distribution of the population.
  - **With noise**: the agents update continuously their choices along the path.

Main Contributions

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
    - **Without noise**: an agent picks a destination prior start moving based on its initial condition and the initial distribution of the population.
    - **With noise**: the agents update continuously their choices along the path.
- Can one anticipate the macroscopic behavior of the population, i.e. the distribution of the agents between the alternatives?

Main Contributions

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
  - **Without noise**: an agent picks a destination prior start moving based on its initial condition and the initial distribution of the population.
  - **With noise**: the agents update continuously their choices along the path.
- Can one anticipate the macroscopic behavior of the population, i.e. the distribution of the agents between the alternatives?
  - We show the existence of at least one infinite population Nash equilibrium.

Main Contributions

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
  - **Without noise**: an agent picks a destination prior start moving based on its initial condition and the initial distribution of the population.
  - **With noise**: the agents update continuously their choices along the path.
- Can one anticipate the macroscopic behavior of the population, i.e. the distribution of the agents between the alternatives?
  - We show the existence of at least one infinite population Nash equilibrium.
  - There exists a one to one map between the infinite population equilibria and the fixed points of a finite dimensional operator F defined on $\mathbb{R}^l$.

Main Contributions

- How do the agents act on the individual level?
  Do they make their choices a priori or continuously while moving ?
  - **Without noise**: an agent picks a destination prior start moving based on its initial condition and the initial distribution of the population.
  - **With noise**: the agents update continuously their choices along the path.
- Can one anticipate the macroscopic behavior of the population, i.e. the distribution of the agents between the alternatives?
  - We show the existence of at least one infinite population Nash equilibrium.
  - There exists a one to one map between the infinite population equilibria and the fixed points of a finite dimensional operator F defined on $\mathbb{R}^l$.
  - **The fixed points of F are the potential distributions of the choices between the alternatives**

    $F(\lambda_1, \ldots, \lambda_l) = (\lambda_1, \ldots, \lambda_l) \Rightarrow \lambda_j\%$ of the agents are closer to p_j at T.

Main Contributions (Cont.)

- What is the necessary information to compute the individual strategies and anticipate the macroscopic behaviors?

Main Contributions (Cont.)

- What is the necessary information to compute the individual strategies and anticipate the macroscopic behaviors?
    - To compute its strategy, each agent needs to know its state and the initial probability distribution of the population **(decentralized strategies)**.

Main Contributions (Cont.)

- What is the necessary information to compute the individual strategies and anticipate the macroscopic behaviors?
  - To compute its strategy, each agent needs to know its state and the initial probability distribution of the population **(decentralized strategies)**.
  - To anticipate the macroscopic behavior, one need to know the initial probability distribution of the population.

Generic Agent's Best Response

MFG approach: $\bar{x}$ is assumed equal to a continuous path $\hat{x}$ and a generic agent minimizes

$$J(u, \hat{x}, x^0) = \int_0^T \left\{ \frac{q}{2}\|x - \hat{x}\|^2 + \frac{r}{2}\|u\|^2 \right\} \mathrm{d}t + \frac{M}{2} \min_{j=1,\dots,l} \left(\|x(T) - p_j\|^2 \right)$$

$$= \min_{j=1,\dots,l} J_j(u, \hat{x}, x^0)$$

$=$ **Minimum of l linear tracking problems, each associated with one of the destination points.**

$$J_j(u, \hat{x}, x^0) = \int_0^T \left\{ \frac{q}{2}\|x - \hat{x}\|^2 + \frac{r}{2}\|u\|^2 \right\} \mathrm{d}t + \frac{M}{2}\|x(T) - p_j\|^2.$$

Recall (Linear Tracking Problem):

- Optimal cost $J_j^*(\hat{x}, x^0)$ = Quadratic function of x^0
- Optimal controller u_j^* = Linear feedback + Tracking term

Generic Agent initially at x^0 → $J_k^*(\hat{x}, x^0) = \min_{j=1,\dots,l} J_j^*(\hat{x}, x^0)$ → Go towards p_k

Basin of attraction:

$$D_j(\hat{x}) = \left\{ x^0 \in \mathbb{R}^n \middle| J_j^*(\hat{x}, x^0) \textbf{ least costly } \right\}$$

$$= \textbf{Region delimited by } l - 1 \textbf{ hyperplans}$$

Generic Agent's Best Response (cont.)

Theorem (Generic Agent's Best Response to $\hat{x}$)

The tracking problem has a unique optimal control law

$$u^* = u_j^* = \text{Optimal control law of } J_j^*(\hat{x}, x^0) \qquad \text{if } x^0 \in D_j(\hat{x})$$

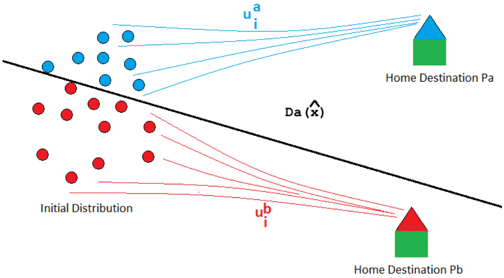

Fixed Point $\bar{x}$

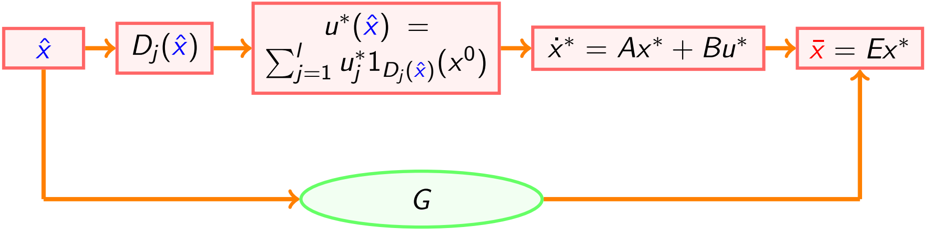

Where are the fixed points of G?

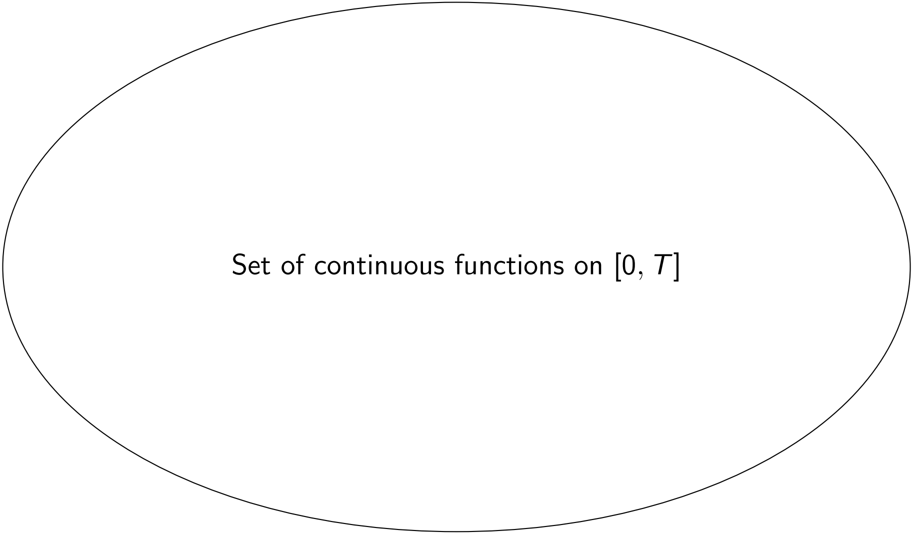

Where are the fixed points of G?

For all $\lambda = (\lambda_1, \ldots, \lambda_l) \in [0,1]^l$,
set of optimal states $y^*(t, \lambda)$ of

$$\int_0^T \frac{r}{2}\|v\|^2 dt + \frac{M}{2}\|y(T) - \sum_{j=1}^{l} \lambda_j p_j\|^2$$

$$\text{s.t. } \dot{y} = Ay + Bv, \quad y(0) = \mu_0$$

Where are the fixed points of G?

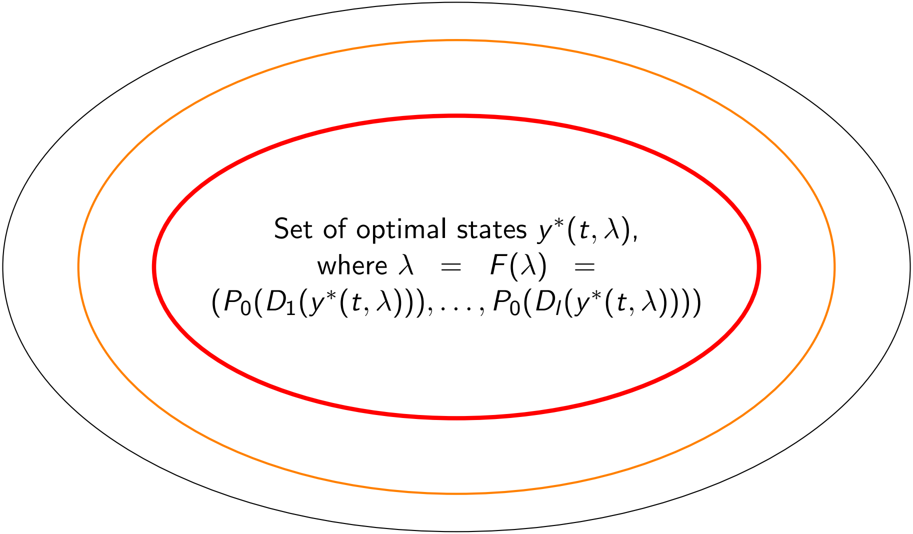

Deterministic Dynamics - Main Theorem

Theorem

The following statements hold:

1. *The set of fixed point trajectories* $\bar{x}(t)$ *is defined by* $\bar{x}(t) = y^*(t, \lambda)$ *for all fixed points* λ *of the finite dimensional function* F *map.*
2. F *has at least one fixed point (equivalently* G *has at least one fixed point).*
3. *The optimal strategy profile* (u_i^*, u_{-i}^*)*, when tracking any of the fixed points* $\bar{x}$*, constitute an* $\epsilon-$*Nash equilibirium, with* ϵ *converges to zero as* N *goes to infinity.*

Proof sketch

- Point 1: Brouwer's fixed point theorem.

Deterministic Dynamics - Summary

- The fixed points paths are totally determined by the fixed points of F.
- The fixed points of F are the potential distributions of the choices between the alternatives. If λ is a fixed point of F then $100\lambda_j\%$ of the agents will go towards p_j.
- Multiple fixed point paths $\bar{x}$ may exist, and are computed as follows

Compute a fixed point λ of F (Broyden's method for example)

↓

Compute the corresponding fixed point path $\bar{x}(t) = y^*(t, \lambda)$

Generic Agent's Best Response - Min-LQG problem

We assume $\bar{x}$ given and call it $\hat{x}$.
The "min-LQG" optimal tracking problem (For clarity purposes, only the binary choice scalar case):

$$J\left(x(0), u(.)\right) = \mathbb{E}\left[\int_0^T \left\{\frac{q}{2}(x-\hat{x})^2 + \frac{r}{2}u^2\right\} dt + \frac{M}{2}\min_{j=1,2}(x(T)-p_j)^2\right]$$

s.t. $dx(t) = \left(ax(t) + bu(t)\right)dt + \sigma dw(t).$

The associated HJB equation:

$$-\frac{\partial V}{\partial t} = ax\frac{\partial V}{\partial x} - \frac{b^2}{2r}\left(\frac{\partial V}{\partial x}\right)^2 + \frac{\sigma^2}{2}\frac{\partial^2 V}{\partial x^2} + \frac{q}{2}(x-\hat{x})^2$$

$$V(T,x) = \frac{M}{2}\min_{j=1,2}(x-p_j)^2.$$

Generic Agent's Best Response - Min-LQG problem (cont.)

Theorem (min-LQG HJB Solution)

The HJB equation of the "min-LQG" problem has a unique solution:

$$V(t,x) = -\frac{\sigma^2 r}{b^2} \log \Bigg(\exp\left(-\frac{b^2}{\sigma^2 r} V_1(t,x)\right) \mathbb{P}\left(x_*^{(1)}(T) \leq c \,\middle|\, x_*^{(1)}(t) = x\right) + \exp\left(-\frac{b^2}{\sigma^2 r} V_2(t,x)\right) \mathbb{P}\left(x_*^{(2)}(T) \geq c \,\middle|\, x_*^{(2)}(t) = x\right) \Bigg),$$

where $c = \frac{p_1+p_2}{2}$, V_j *and* $x_*^{(j)}$ *are the optimal cost-to-go and optimal state of the LQG problem*

$$J^{(j)}\left(x^{(j)}(0), u^{(j)}(.)\right) = \mathbb{E}\Big[\int_0^T \left\{ \frac{q}{2}(x^{(j)} - \hat{x})^2 + \frac{r}{2}(u^{(j)})^2 \right\} dt + \frac{M}{2}(x^{(j)}(T) - p_j)^2 \Big]$$

s.t. the dynamics of the generic agent.

Generic Agent's Best Response- Min-LQG problem (cont.)

Theorem (Generic agent's best response)

$$u_*(t,x,\hat{x}) = \sum_{j=1}^{2} \frac{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_j(t,x)\right)}{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_1(t,x)\right) + \exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_2(t,x)\right)} u_*^{(j)}(t,x)$$

$$\tilde{V}_j(t,x) = V_j(t,x) - \frac{\sigma^2 r}{b^2}\log\left(g_j(t,x)\right),$$

where $g_j(t,x) = \mathbb{P}\left(x_*^{(j)}(T) \text{ closer to } p_j \middle| x_*^{(j)}(t) = x\right)$

Generic Agent's Best Response- Min-LQG problem (cont.)

$$u_*(t,x,\hat{x}) = \sum_{j=1}^{2} \frac{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_j(t,x)\right)}{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_1(t,x)\right) + \exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_2(t,x)\right)} u_*^{(j)}(t,x)$$

$$\tilde{V}_j(t,x) = V_j(t,x) - \frac{\sigma^2 r}{b^2}\log\left(g_j(t,x)\right),$$

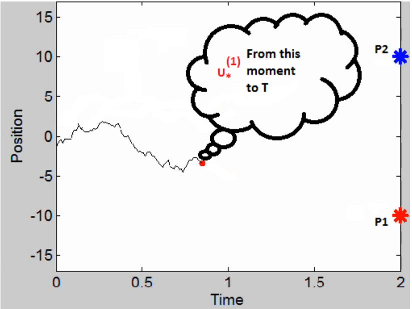

$$u_*(t,x,\hat{x}) = \sum_{j=1}^{2} \frac{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_j(t,x)\right)}{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_1(t,x)\right) + \exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_2(t,x)\right)} u_*^{(j)}(t,x)$$

$$\tilde{V}_1(t,x) = V_1(t,x) - \frac{\sigma^2 r}{b^2}\log\left(g_1(t,x)\right)$$

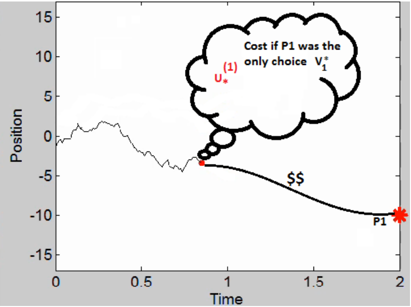

Generic Agent's Best Response- Min-LQG problem (cont.)

$$u_*(t,x,\hat{x}) = \sum_{j=1}^{2} \frac{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_j(t,x)\right)}{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_1(t,x)\right) + \exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_2(t,x)\right)} u_*^{(j)}(t,x)$$

$$\tilde{V}_1(t,x) = V_1(t,x) - \frac{\sigma^2 r}{b^2}\log\left(g_1(t,x)\right)$$

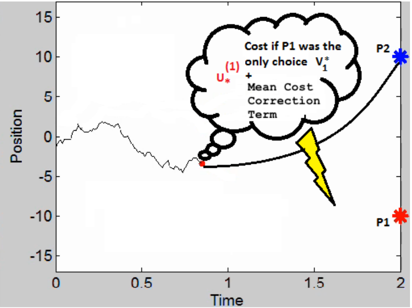

$$u_*(t,x,\hat{x})=\sum_{j=1}^{2}\frac{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_j(t,x)\right)}{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_1(t,x)\right)+\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_2(t,x)\right)}u_*^{(j)}(t,x)$$

$$\tilde{V}_j(t,x) = V_j(t,x) - \frac{\sigma^2 r}{b^2}\log\left(g_j(t,x)\right),$$

$$g_j(t,x) = P\left(x_*^{(j)}(T) \text{ closer to } p_j \middle| x_*^{(j)}(t) = x\right)$$

- The generic agent is indecisive: its best response is a convex combination of $u_*^{(1)}$ (choosing p_1) and $u_*^{(2)}$ (choosing p_2).

Generic Agent's Best Response- Min-LQG problem (cont.)

$$u_*(t,x,\hat{x})=\sum_{j=1}^{2}\frac{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_j(t,x)\right)}{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_1(t,x)\right)+\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_2(t,x)\right)}u_*^{(j)}(t,x)$$

$$\tilde{V}_j(t,x) = V_j(t,x) - \frac{\sigma^2 r}{b^2}\log\left(g_j(t,x)\right),$$

$$g_j(t,x) = P\left(x_*^{(j)}(T) \text{ closer to } p_j \middle| x_*^{(j)}(t) = x\right)$$

- The generic agent is indecisive: its best response is a convex combination of $u_*^{(1)}$ (choosing p_1) and $u_*^{(2)}$ (choosing p_2).
- $u_*^{(1)}$ and $u_*^{(2)}$ are weighted by a Gibbs distribution, which attributes more mass to the less costly and risky choice.

Fixed Point $\bar{x}$

A fixed point path $\bar{x}$ satisfies the following Mean Field equations:

$$\begin{aligned} dx_*(t) &= (ax_*(t) + bu_*(t, x_*(t), \bar{x}))\, dt + \sigma dw(t), \quad x_*(0) = x(0) \\ \bar{x}(t) &= \mathbb{E}x_*(t). \end{aligned} \tag{3}$$

(3) is a nonlinear Mckean-Vlasov equation.

- Show that the set of solutions of (3) can be one to one mapped to the set of fixed points of a finite dimensional operator.
- At least one solution of (3) exists.
- As in the deterministic case, the fixed points of the finite dimensional operator are the distributions of the choices.

Where are the fixed points of paths?

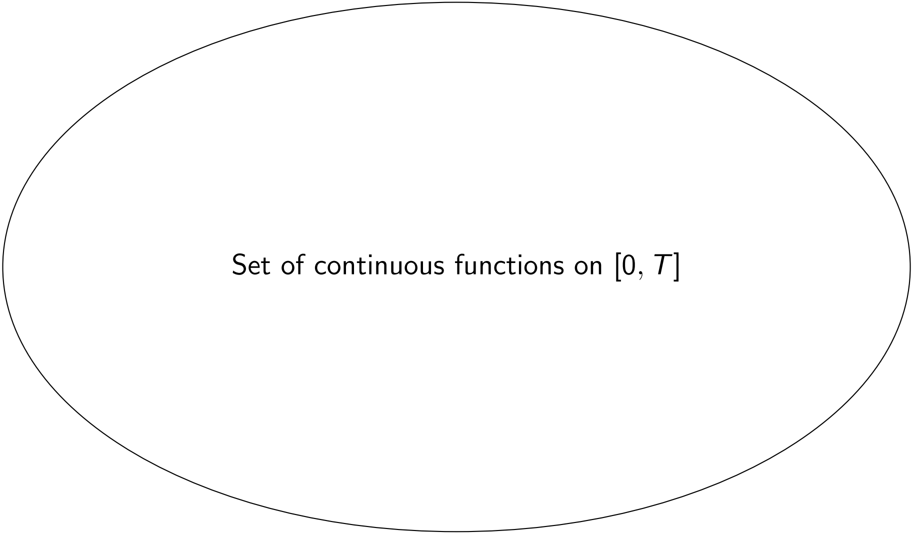

Where are the fixed points of paths?

For all $\lambda \in [0,1]$, set of optimal states $y^*(t,\lambda)$ of

$$\int_0^T \frac{r}{2}\|v\|^2 dt + \frac{M}{2}\|y(T) - \lambda p_1 - (1-\lambda)p_2\|^2$$

$$\text{s.t. } \dot{y} = Ay + Bv, \quad y(0) = \mu_0$$

Where are the fixed points of paths?

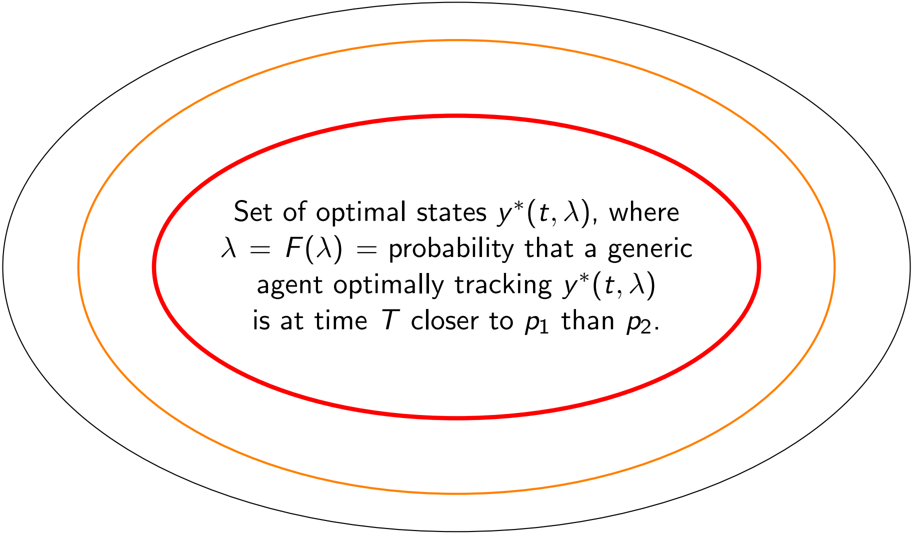

Fixed Point $\bar{x}$ (Cont.)

Theorem

The following statements hold:

1. *The set of fixed point trajectories $\bar{x}(t)$ is defined by $\bar{x}(t) = y^*(t, \lambda)$ for all fixed points λ of the F map.*
2. *F has at least one fixed point (equivalently G has at least one fixed point).*
3. *The optimal strategy profile (u_i^*, u_{-i}^*), when tracking any of the fixed points $\bar{x}$, constitute an $\epsilon-$Nash equilibirium, with ϵ converges to zero as N goes to infinity.*

Proof sketch

- Point 1: Brouwer's fixed point theorem.

Fixed Point $\bar{x}$ (Cont.)

If λ is a fixed point of F then $100\lambda\%$ of the agents will be closer to p_1 at time T.

Computation of a fixed point of F - Bisection Method

$\lambda_0 = 0, \lambda_1 = 1$

$\lambda = \frac{\lambda_0+\lambda_1}{2}, \; e = |\lambda_0 - \lambda_1|$ ← While $e > e_d$

Implicit finite difference Scheme for the Fokker-Planck equation of

$$dx_*^\lambda(t) = \left(ax_*^\lambda(t) + bu_*\left(t, x_*^\lambda(t), y^*(t,\lambda)\right)\right) dt + \sigma dw(t)$$

Compute $F(\lambda) = P(x_*^\lambda(T) \leq c)$

If $(F(\lambda)-\lambda)(F(\lambda_0)-\lambda_0) < 0$, $\lambda_1 = \lambda$. Else $\lambda_0 = \lambda$.

Simulations

Influence of the social effect (q) on the number of equilibria and the size of the groups.

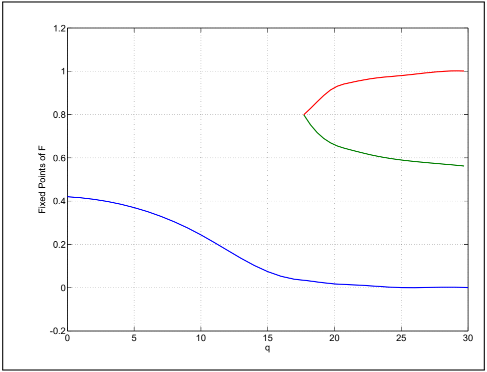

Deterministic Vs. Stochastic dynamics

Deterministic dynamics

- Decisive Agents: $u^* = 1_{D_1(\bar{x})}(x^0)u_*^{(1)} + (1 - 1_{D_1(\bar{x})}(x^0))u_*^{(2)}$.

Stochastic dynamics

- Indecisive Agents: $u^* = p(t, x, \bar{x})u_*^{(1)} + (1 - p(t, x, \bar{x}))u_*^{(2)}$.

Deterministic Vs. Stochastic dynamics

Deterministic dynamics

- Decisive Agents: $u^* = 1_{D_1(\bar{x})}(x^0)u_*^{(1)} + (1 - 1_{D_1(\bar{x})}(x^0))u_*^{(2)}$.
- $\bar{x}$ is a fixed point if and only if $\bar{x}(t) = y^*(t, \lambda)$, where λ is a fixed point of a well defined finite dimensional operator F.

Stochastic dynamics

- Indecisive Agents: $u^* = p(t, x, \bar{x})u_*^{(1)} + (1 - p(t, x, \bar{x}))u_*^{(2)}$.
- $\bar{x}$ is a fixed point if and only if $\bar{x}(t) = y^*(t, \lambda)$, where λ is a fixed point of a well defined finite dimensional operator F.

Deterministic Vs. Stochastic dynamics

Deterministic dynamics

- Decisive Agents: $u^* = 1_{D_1(\bar{x})}(x^0)u_*^{(1)} + (1 - 1_{D_1(\bar{x})}(x^0))u_*^{(2)}$.
- $\bar{x}$ is a fixed point if and only if $\bar{x}(t) = y^*(t, \lambda)$, where λ is a fixed point of a well defined finite dimensional operator F.
- $Px(0)^{-1}$ and $\lambda \to F(\lambda)$.

Stochastic dynamics

- Indecisive Agents: $u^* = p(t, x, \bar{x})u_*^{(1)} + (1 - p(t, x, \bar{x}))u_*^{(2)}$.
- $\bar{x}$ is a fixed point if and only if $\bar{x}(t) = y^*(t, \lambda)$, where λ is a fixed point of a well defined finite dimensional operator F.
- $Px(0)^{-1}$ and $\lambda \to Px_*^{\lambda}(T)^{-1}$ and $\lambda \to F(\lambda)$, where

$$dx_*^{\lambda}(t) = ax_*^{\lambda}(t)dt + bu_*\Big(t, x_*^{\lambda}(t), y^*(t, \lambda)\Big)dt + \sigma dw(t)$$

“Min-LQG” Vs. Static Discrete Choice Models

Economics - Static Discrete Choice Models [McFadden ’74]

Agent

“Min-LQG” Vs. Static Discrete Choice Models

Economics - Static Discrete Choice Models [McFadden ’74]

"Min-LQG" Vs. Static Discrete Choice Models

Economics - Static Discrete Choice Models [McFadden '74]

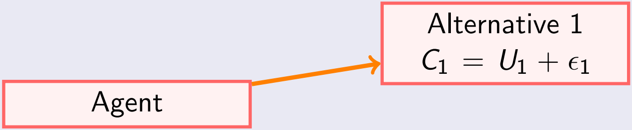

"Min-LQG" Vs. Static Discrete Choice Models

Economics - Static Discrete Choice Models [McFadden '74]

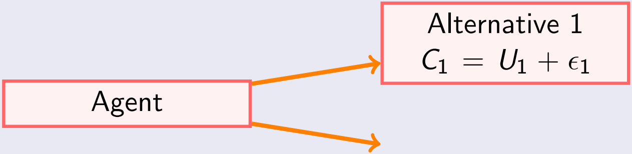

"Min-LQG" Vs. Static Discrete Choice Models

Economics - Static Discrete Choice Models [McFadden '74]

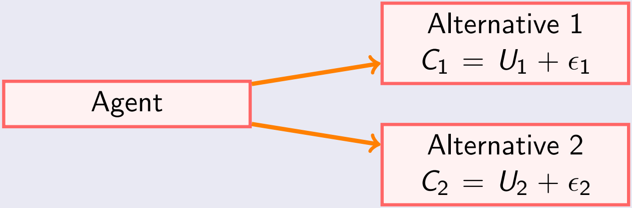

“Min-LQG” Vs. Static Discrete Choice Models

Economics - Static Discrete Choice Models [McFadden ’74]

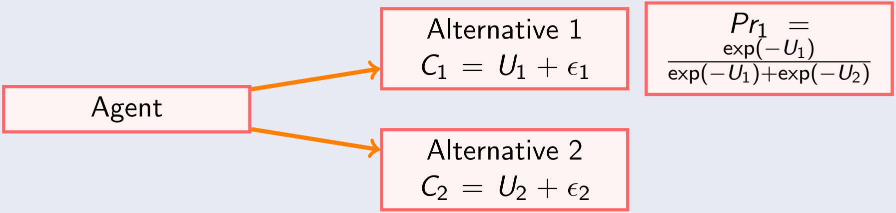

"Min-LQG" Vs. Static Discrete Choice Models

Economics - Static Discrete Choice Models [McFadden '74]

Agent

Alternative 1: $C_1 = U_1 + \epsilon_1$

$$Pr_1 = \frac{\exp(-U_1)}{\exp(-U_1)+\exp(-U_2)}$$

Alternative 2: $C_2 = U_2 + \epsilon_2$

$$Pr_2 = \frac{\exp(-U_2)}{\exp(-U_1)+\exp(-U_2)}$$

"Min-LQG" Vs. Static Discrete Choice Models

Economics - Static Discrete Choice Models [McFadden '74]

Agent

Alternative 1

$C_1 = U_1 + \epsilon_1$

$Pr_1 = \frac{\exp(-U_1)}{\exp(-U_1)+\exp(-U_2)}$

Alternative 2

$C_2 = U_2 + \epsilon_2$

$Pr_2 = \frac{\exp(-U_2)}{\exp(-U_1)+\exp(-U_2)}$

Min-LQG

$$u_*(t,x) = \sum_{j=1}^{2} \frac{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_j(t,x)\right)}{\exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_1(t,x)\right) + \exp\left(-\frac{b^2}{\sigma^2 r}\tilde{V}_2(t,x)\right)} u_*^{(j)}(t,x)$$

$$\tilde{V}_j(t,x) = V_j(t,x) - \frac{\sigma^2 r}{b^2} \log\left(P\left(x_*^{(j)}(T) \text{ closer to } p_j \middle| x_*^{(j)}(t) = x\right)\right)$$

Conclusion

1. We considered a dynamic collective choice problem based on a Min-LQG dynamic game.

2. We showed that multiple $\epsilon-$Nash equilibria may exist. Each equilibrium is characterized by a vector λ describing the way the population splits under a social effect between the destination points.

3. In the deterministic case, an agent picks its destination point prior start moving.

4. In the stochastic case, the agents are indecisive.

Extensions

Dynamic collective choice with nonuniform costs and dynamics, and initial preferences towards the destination points

Dynamic collective choice with an advertiser [CDC'16]

A Stackelberg competition involving:

- A group of agents choosing between two alternatives under the social and advertisement effects.
- An advertiser making some investments to advertise to one of the alternatives.

Dynamic collective choice: social optima

A cooperative game, where a group of agents are cooperatively choosing among a set of alternatives under the social effect.

Thank you