Optimal incentives to mitigate epidemics: A Stackelberg mean field game approach

Alexander Aurell

ORFE, Princeton University

December 7, 2020

Joint work with René Carmona, Gökçe Dayanikli, and Mathieu Laurière (ORFE)

Introduction 1/2

In the absence of a vaccine, how to **incentivize** the individuals of society to make the right **effort** in the fight against an epidemic?

Introduction 1/2

In the absence of a vaccine, how to **incentivize** the individuals of society to make the right **effort** in the fight against an epidemic?

A policy maker's problem: give incentives and penalties to the population that

1. the populations accepts and follows
2. yields a behavior that "controls" the epidemic

How can we encourage risk-averse behavior and reward it optimally?

Introduction 2/2

This talk is based on the approach explored in *"Optimal incentives to mitigate epidemics: A Stackelberg mean field game approach"* A., Carmona, Dayanikli, Lauriére, arXiv 2020.

- The society consists of one **principal** and a **large population** of **agents**.
- How the disease spreads depends on the **agents' efforts** to **slow spread**.
- The agents are not cooperating! They are playing a **mean field game**.
- Principal **optimizes** a contract given knowledge of the agents' response.

Introduction 2/2

This talk is based on the approach explored in *"Optimal incentives to mitigate epidemics: A Stackelberg mean field game approach"* A., Carmona, Dayanikli, Lauriére, arXiv 2020.

- The society consists of one **principal** and a **large population** of **agents**.
- How the disease spreads depends on the **agents' efforts** to **slow spread**.
- The agents are not cooperating! They are playing a **mean field game**.
- Principal **optimizes** a contract given knowledge of the agents' response.

The principal and the population play a **Stackelberg game**.

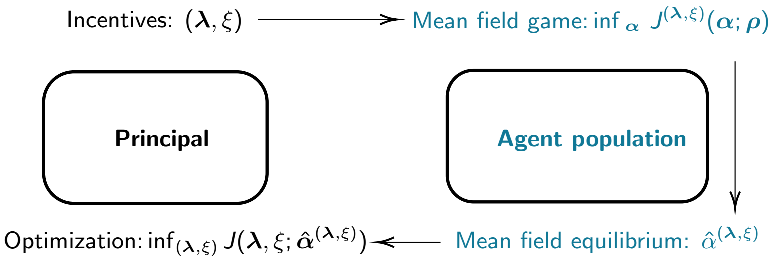

Compartmental models of epidemics 1/4

Epidemic modelling with the SIR model

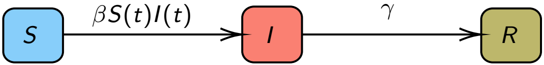

Individuals are categorized either as "**S**usceptible", "**I**nfected" or "**R**emoved".

Compartmental models of epidemics 1/4

Epidemic modelling with the SIR model

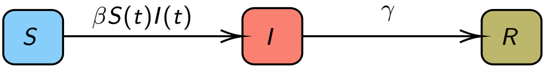

Individuals are categorized either as "**S**usceptible", "**I**nfected" or "**R**emoved".

The system of equation that describes the evolution of the epidemic:

$$\begin{cases} \dot{S}(t) = -\beta S(t)I(t), \quad S(0) \geq 0 \\ \dot{I}(t) = \beta S(t)I(t) - \gamma I(t), \quad I(0) \geq 0 \\ \dot{R}(t) = \gamma I(t), \quad R(0) \geq 0 \\ S(0) + I(0) + R(0) = 1, \end{cases}$$

Many, many variations!

Compartmental models of epidemics 2/4

The epidemic's dynamics is described by two parameters: β and γ.

- **Recovery rate** γ, the reciprocal average infectious time.
- **Transmission rate** β.

What is a reasonable model for agent control of the transmission rate?

Compartmental models of epidemics 2/4

The epidemic's dynamics is described by two parameters: β and γ.

- **Recovery rate** γ, the reciprocal average infectious time.
- **Transmission rate** β.

What is a reasonable model for agent control of the transmission rate?

- In a meeting, does the risk of infection depend on all the meeting parties effort to reduce the transmission rate? Linearly or non-linearly?

Compartmental models of epidemics 2/4

The epidemic's dynamics is described by two parameters: β and γ.

- **Recovery rate** γ, the reciprocal average infectious time.
- **Transmission rate** β.

What is a reasonable model for agent control of the transmission rate?

- In a meeting, does the risk of infection depend on all the meeting parties effort to reduce the transmission rate? Linearly or non-linearly?

- Should effort to reduce transmission rate be universal or state-dependent? Lock down only for the sick or for all?

We argue that β, if controlled, can depend on the action of **many agents**...

Compartmental models of epidemics 3/4

Consider N agents. Agent $i \in \{1, \ldots, N\}$ has state $X_t^i \in \{S, I, R\}$ at time t.

- Meetings in the population occur pairwise and at random with rate β.

Compartmental models of epidemics 3/4

Consider N agents. Agent $i \in \{1, \dots, N\}$ has state $X_t^i \in \{S, I, R\}$ at time t.

- Meetings in the population occur pairwise and at random with rate β.
- If a susceptible agent meets an infected agent, she is infected.

Compartmental models of epidemics 3/4

Consider N agents. Agent $i \in \{1, \dots, N\}$ has state $X_t^i \in \{S, I, R\}$ at time t.

- Meetings in the population occur pairwise and at random with rate β.
- If a susceptible agent meets an infected agent, she is infected.
- The recovery rate is γ.

The population of agents is described by an interacting system of (continuous time) exchangeable **Markov chains** with transition rate matrix

$$Q(p_t^N) = \begin{bmatrix} -\beta p_t^N(I) & \beta p_t^N(I) & 0 \\ 0 & -\gamma & \gamma \\ 0 & 0 & 0 \end{bmatrix}$$

where $p_t^N(I)$ is the **proportion** of the population that is infected at time t,

$$p_t^N = (p_t^N(S), p_t^N(I), p_t^N(R)) := \left(\frac{1}{N} \sum_{j=1}^{N} \mathbb{1}_i(X_t^j) \right)_{i \in \{S,I,R\}}$$

Compartmental models of epidemics 4/4

What if the agents can take precautions so that a meeting does not automatically lead to infection?

- The probability of infection is decreased by the action/effort of two agents that meet in a multiplicative way.

The agents control their "contact factor".

Compartmental models of epidemics 4/4

What if the agents can take precautions so that a meeting does not automatically lead to infection?

- The probability of infection is decreased by the action/effort of two agents that meet in a multiplicative way.

The agents control their "contact factor".

With **contact factor** control, agent j's transition rate from S to I:

$$\beta\alpha_t^j \frac{1}{N}\sum_{k=1}^{N}\alpha_t^k \mathbb{1}_I(X_{t-}^k)$$

- equals the SIR rate $\beta p_t^N(I)$ if $\alpha_t^j = 1$, $j = 1, \ldots, N$.

Compartmental models of epidemics 4/4

What if the agents can take precautions so that a meeting does not automatically lead to infection?

- The probability of infection is decreased by the action/effort of two agents that meet in a multiplicative way.

The agents control their "contact factor".

With **contact factor** control, agent j's transition rate from S to I:

$$\beta \alpha_t^j \frac{1}{N} \sum_{k=1}^{N} \alpha_t^k \mathbb{1}_I(X_{t-}^k)$$

- equals the SIR rate $\beta p_t^N(I)$ if $\alpha_t^j = 1$, $j = 1, \ldots, N$.

Symmetric, weak interaction ... MFG?

Mean field games 1/3

Idea from statistical physics:

- N players in a game
- Interactions between players' states
  - in the coefficients of the state dynamics
  - in the cost functions
- exclusively through the **empirical distribution**

$$\mu_t^N = \frac{1}{N} \sum_{j=1}^{N} \delta_{X_t^j}$$

Mean field games 1/3

Idea from statistical physics:

- N players in a game
- Interactions between players' states
  - in the coefficients of the state dynamics
  - in the cost functions
- exclusively through the **empirical distribution**

$$\mu_t^N = \frac{1}{N} \sum_{j=1}^{N} \delta_{X_t^j}$$

Consequences:

- Strong symmetry among the players
- Each player can hardly influence the system when N is large.

Mean field games 1/3

Idea from statistical physics:

- N players in a game
- Interactions between players' states
  - in the coefficients of the state dynamics
  - in the cost functions
- exclusively through the **empirical distribution**

$$\mu_t^N = \frac{1}{N}\sum_{j=1}^{N}\delta_{X_t^j}$$

Consequences:

- Strong symmetry among the players
- Each player can hardly influence the system when N is large.

Mean field game (MFG): the limit game as $N \to \infty$

$$(i)\ \hat{\alpha} = \arg\inf_{\alpha} J(\alpha;\hat{\mu}), \quad (ii)\ \hat{\mu} = \text{distribution of } X^{\hat{\alpha}}$$

Lasry-Lions (2006), Huang-Malhamé-Caines (2006)

Mean field games 2/3

$$\beta \alpha_t^j \frac{1}{N} \sum_{k=1}^{N} \alpha_t^k \mathbb{1}_I(X_{t-}^k)$$

We anticipate that, for very large N, we can approximate the game with **contact factor** control with an **extended finite-state MFG**.

Mean field games 2/3

$$\beta\alpha_t^j \frac{1}{N}\sum_{k=1}^{N}\alpha_t^k \mathbb{1}_I(X_{t-}^k)$$

We anticipate that, for very large N, we can approximate the game with **contact factor** control with an **extended finite-state MFG**.

Transition rate matrix

$$Q(t,\alpha,\rho) = \begin{bmatrix} -\beta\alpha_t \int_A a\rho_t(da, I) & \beta\alpha_t \int_A a\rho_t(da, I) & 0 \\ 0 & -\gamma & \gamma \\ 0 & 0 & \cdots \end{bmatrix},$$

where ρ_t is a **joint state-and-control distribution**.

Gomes *et al* (2010, 2013), Kolokoltsov (2012), Carmona-Wang (2016, 2018), Cecchin-Fischer (2018), Bayraktar-Cohen (2018), Choutri *et al* (2018, 2019).

Mean field games 3/3

Motivated by the SIR example, we will consider a MFG with:

- finite state space
- **extended** mean field interaction, i.e., interaction through the joint state-control distribution ρ

for the purpose of modeling decision making during an epidemic.

Elie *et al* (2020), Hubert *et al* (2020), Charpentier *et al* (2020), Cho (2020)

Motivated by the SIR example, we will consider a MFG with:

- finite state space
- **extended** mean field interaction, i.e., interaction through the joint state-control distribution ρ

for the purpose of modeling decision making during an epidemic.

Elie *et al* (2020), Hubert *et al* (2020), Charpentier *et al* (2020), Cho (2020)

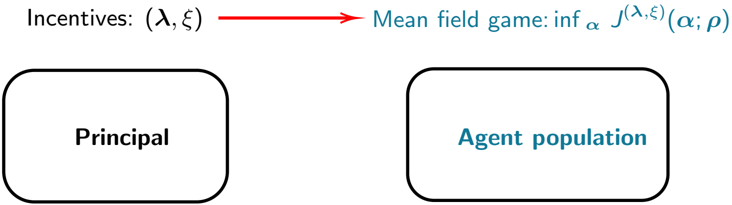

But first, some notation ...

MFG for epidemics with contract factor control 1/3

Setup

- **Sample space** Ω càdlàg functions $\omega : [0, T] \to E := \{e_1, \dots, e_m\}$
- **Canonical process** $\boldsymbol{X}$: $X_t(\omega) = \omega(t)$.
- **Filtration** $\mathbb{F}$ natural filtration generated by $\boldsymbol{X}$ and $\mathcal{F} := \mathcal{F}_T$.

MFG for epidemics with contract factor control 1/3

Setup

- **Sample space** Ω càdlàg functions $\omega : [0, T] \to E := \{e_1, \dots, e_m\}$
- **Canonical process** $\boldsymbol{X}$: $X_t(\omega) = \omega(t)$.
- **Filtration** $\mathbb{F}$ natural filtration generated by $\boldsymbol{X}$ and $\mathcal{F} := \mathcal{F}_T$.
- **Basic transition rate matrix** Q^0: rate from e_i to e_j equal to 1 if $(i, j) \in G \subset \{1, \dots, m\}^2$, otherwise zero.

MFG for epidemics with contract factor control 1/3

Setup

- **Sample space** Ω càdlàg functions $\omega : [0, T] \to E := \{e_1, \dots, e_m\}$
- **Canonical process** $\boldsymbol{X}$: $X_t(\omega) = \omega(t)$.
- **Filtration** $\mathbb{F}$ natural filtration generated by $\boldsymbol{X}$ and $\mathcal{F} := \mathcal{F}_T$.
- **Basic transition rate matrix** Q^0: rate from e_i to e_j equal to 1 if $(i,j) \in G \subset \{1, \dots, m\}^2$, otherwise zero.
- **Basic probability space** $(\Omega, \mathbb{F}, \mathcal{F}, \mathbb{P})$ such that
  - $\mathbb{P} \circ X_0^{-1} = p^0 \in \mathcal{P}(E)$
  - $\boldsymbol{X}$ Markov chain with transition rate matrix Q^0
- Under $\mathbb{P}$ $\boldsymbol{X}$ has the representation

$$X_t = X_0 + \int_0^t X_{s-}^* Q^0 ds + \mathcal{M}_t \tag{1}$$

MFG for epidemics with contract factor control 2/3

Controlled probability space

- **Control processes** $\mathbb{A}$ A-valued $\mathbb{F}$-predictable processes and $A := [0, 1]$.
- **Action-state laws** $\mathcal{R} := \mathcal{P}(A \times E)$ Borel probability measures on $A \times E$.
- **Measure flows** $M(\mathcal{R})$ and $M(\mathcal{P}(E))$ measurable mappings from $[0, T]$ to $\mathcal{R}$ and $\mathcal{P}(E)$, respectively.
- **Metrics:** A Euclidean metric, E bounded discrete metric, $A \times E$ 1-product metric, $\mathcal{P}(E)$ Euclidean metric (on the simplex), $\mathcal{R}$ 1-Wasserstein metric W_R.

MFG for epidemics with contract factor control 2/3

Controlled probability space

- **Control processes** $\mathbb{A}$ A-valued $\mathbb{F}$-predictable processes and $A := [0, 1]$.
- **Action-state laws** $\mathcal{R} := \mathcal{P}(A \times E)$ Borel probability measures on $A \times E$.
- **Measure flows** $M(\mathcal{R})$ and $M(\mathcal{P}(E))$ measurable mappings from $[0, T]$ to $\mathcal{R}$ and $\mathcal{P}(E)$, respectively.
- **Metrics:** A Euclidean metric, E bounded discrete metric, $A \times E$ 1-product metric, $\mathcal{P}(E)$ Euclidean metric (on the simplex), $\mathcal{R}$ 1-Wasserstein metric W_R.
- **Controlled transition rate matrix** $Q : [0, T] \times A \times \mathcal{R} \to \mathbb{R}^{m \times m}$ bounded, measurable function such that $Q(t, \alpha, \rho)$ is a Q-matrix.

MFG for epidemics with contract factor control 2/3

Controlled probability space

- **Control processes** $\mathbb{A}$ A-valued $\mathbb{F}$-predictable processes and $A := [0, 1]$.
- **Action-state laws** $\mathcal{R} := \mathcal{P}(A \times E)$ Borel probability measures on $A \times E$.
- **Measure flows** $M(\mathcal{R})$ and $M(\mathcal{P}(E))$ measurable mappings from $[0, T]$ to $\mathcal{R}$ and $\mathcal{P}(E)$, respectively.
- **Metrics:** A Euclidean metric, E bounded discrete metric, $A \times E$ 1-product metric, $\mathcal{P}(E)$ Euclidean metric (on the simplex), $\mathcal{R}$ 1-Wasserstein metric W_R.
- **Controlled transition rate matrix** $Q : [0, T] \times A \times \mathcal{R} \to \mathbb{R}^{m \times m}$ bounded, measurable function such that $Q(t, \alpha, \rho)$ is a Q-matrix.
- For $(\boldsymbol{\alpha}, \boldsymbol{\rho}) \in \mathbb{A} \times M(\mathcal{R})$ the probability measure $\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}}$ on $(\Omega, \mathcal{F})$ is given by $d\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}} = \mathcal{E}_T d\mathbb{P}$ on $\mathcal{F}$ where

$$\begin{aligned} \mathcal{E}_t &= 1 + \int_0^t \mathcal{E}_{s-} X_{s-}^* \left(Q(s, \alpha_s, \rho_s) - Q^0 \right) \psi_s^+ d\mathcal{M}_s, \\ \psi_t &:= \mathrm{diag}(Q^0 X_{t-}) - Q^0 \mathrm{diag}(X_{t-}) - \mathrm{diag}(X_{t-}) Q^0 \end{aligned} \tag{2}$$

MFG for epidemics with contract factor control 2/3

Controlled probability space

- **Control processes** $\mathbb{A}$ A-valued $\mathbb{F}$-predictable processes and $A := [0, 1]$.
- **Action-state laws** $\mathcal{R} := \mathcal{P}(A \times E)$ Borel probability measures on $A \times E$.
- **Measure flows** $M(\mathcal{R})$ and $M(\mathcal{P}(E))$ measurable mappings from $[0, T]$ to $\mathcal{R}$ and $\mathcal{P}(E)$, respectively.
- **Metrics:** A Euclidean metric, E bounded discrete metric, $A \times E$ 1-product metric, $\mathcal{P}(E)$ Euclidean metric (on the simplex), $\mathcal{R}$ 1-Wasserstein metric W_R.
- **Controlled transition rate matrix** $Q : [0, T] \times A \times \mathcal{R} \to \mathbb{R}^{m \times m}$ bounded, measurable function such that $Q(t, \alpha, \rho)$ is a Q-matrix.
- For $(\boldsymbol{\alpha}, \boldsymbol{\rho}) \in \mathbb{A} \times M(\mathcal{R})$ the probability measure $\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}}$ on $(\Omega, \mathcal{F})$ is given by $d\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}} = \mathcal{E}_T d\mathbb{P}$ on $\mathcal{F}$ where

$$
\begin{aligned}
\mathcal{E}_t &= 1 + \int_0^t \mathcal{E}_{s-} X_{s-}^* \left(Q(s, \alpha_s, \rho_s) - Q^0 \right) \psi_s^+ \, d\mathcal{M}_s, \\
\psi_t &:= \mathrm{diag}(Q^0 X_{t-}) - Q^0 \mathrm{diag}(X_{t-}) - \mathrm{diag}(X_{t-}) Q^0
\end{aligned}
\tag{2}
$$

- Under $\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}}$, $\boldsymbol{X}$ is a Markov chain with transition rate matrix $Q(t, \alpha_t, \rho_t)$ at time t.

MFG for epidemics with contract factor control 3/3

The agents' problem: find the mean-field Nash equilibrium.

The cost for $\boldsymbol{\alpha} \in \mathbb{A}$ is

$$J^{\boldsymbol{\lambda},\xi}(\boldsymbol{\alpha};\boldsymbol{\rho}) := \mathbb{E}^{\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}}}\left[\int_0^T f(t, X_t, \alpha_t, \rho_t; \lambda_t)dt - U(\xi)\right],$$

where

$(\boldsymbol{\lambda}, \xi)$	principal's policy choice, the **contract**
$f : [0, T] \times E \times A \times \mathcal{R} \to \mathbb{R}$	running cost, depends on $\boldsymbol{\lambda}$
$U : \mathbb{R} \to \mathbb{R}$	utility of a terminal payment
$\boldsymbol{\rho} = (\rho_t)_{t\in[0,T]} \in M(\mathcal{R})$	joint state-control distribution in the population
$\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}} \in \mathcal{P}(\Omega, \mathcal{F})$	under which X_t has rate matrix $Q(t, \alpha_t, \rho_t)$

MFG for epidemics with contract factor control 3/3

The agents' problem: find the mean-field Nash equilibrium.

The cost for $\boldsymbol{\alpha} \in \mathbb{A}$ is

$$J^{\boldsymbol{\lambda},\xi}(\boldsymbol{\alpha};\boldsymbol{\rho}) := \mathbb{E}^{\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}}}\left[\int_0^T f(t, X_t, \alpha_t, \rho_t; \lambda_t)dt - U(\xi)\right],$$

where

$(\boldsymbol{\lambda}, \xi)$	principal's policy choice, the **contract**
$f : [0,T] \times E \times A \times \mathcal{R} \to \mathbb{R}$	running cost, depends on $\boldsymbol{\lambda}$
$U : \mathbb{R} \to \mathbb{R}$	utility of a terminal payment
$\boldsymbol{\rho} = (\rho_t)_{t\in[0,T]} \in M(\mathcal{R})$	joint state-control distribution in the population
$\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{\rho}} \in \mathcal{P}(\Omega, \mathcal{F})$	under which X_t has rate matrix $Q(t, \alpha_t, \rho_t)$

Definition

If the pair $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}) \in \mathbb{A} \times M(\mathcal{R})$ satisfies

(i) $\hat{\boldsymbol{\alpha}} = \arg\inf_{\boldsymbol{\alpha}\in\mathbb{A}} J^{\boldsymbol{\lambda},\xi}(\boldsymbol{\alpha}, \hat{\boldsymbol{\rho}})$;

(ii) $\forall t \in [0,T] \ : \ \hat{\rho}_t = \mathbb{Q}^{\hat{\boldsymbol{\alpha}},\hat{\boldsymbol{\rho}}} \circ (\hat{\alpha}_t, X_t)^{-1}$,

then $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}})$ is a **mean-field Nash equilibrium** given the contract $(\boldsymbol{\lambda}, \xi)$.

Characterizing mean-field Nash equilibria 1/3

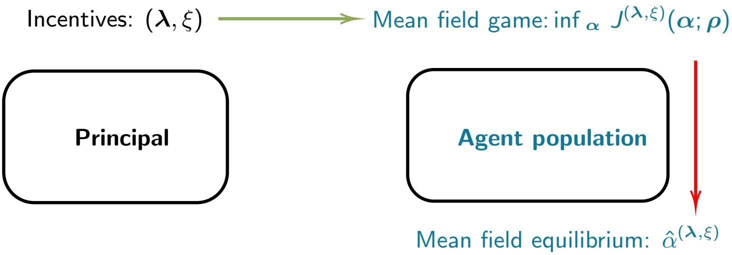

$\mathcal{N}(\boldsymbol{\lambda}, \xi) :=$ the set of mean field Nash equilibria given the contract $(\boldsymbol{\lambda}, \xi)$.

A **forward-backward SDE** (FBSDE) helps us solving for $\mathcal{N}(\boldsymbol{\lambda}, \xi)$...

Characterizing mean-field Nash equilibria 2/3

Under suitable assumptions $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}) \in \mathcal{N}(\boldsymbol{\lambda}, \xi)$ if $(\boldsymbol{Y}, \boldsymbol{Z}, \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}, \mathbb{Q})$ solves [1] the FBSDE

$$\begin{cases} Y_t = U(\xi) + \int_t^T \hat{H}(s, X_{s-}, Z_s, \hat{\rho}_s) ds - \int_t^T Z_s^* d\mathcal{M}_s, \\ \mathcal{E}_t = 1 + \int_0^t \mathcal{E}_{s-} X_{s-}^* \left(Q(s, \hat{\alpha}_s, \hat{\rho}_s) - Q^0 \right) \psi_s^+ d\mathcal{M}_s, \\ \hat{\rho}_t = \mathbb{Q} \circ (\hat{\alpha}_t, X_t)^{-1}, \quad \frac{d\mathbb{Q}}{d\mathbb{P}} = \mathcal{E}_T, \quad \hat{\alpha}_t = \hat{a}(t, X_{t-}, Z_t, \hat{\rho}_t), \end{cases} \tag{3}$$

where $\hat{H}$ is the minimized Hamiltonian and $\mathcal{M}$ is the canonical process' compensating martingale (under $\mathbb{P}$):

- $H : (t, x, z, \alpha, \rho) \mapsto x^* \left(Q(t, \alpha, \rho) - Q^0 \right) z + f(t, x, \alpha, \rho; \lambda_t)$
- $X_t = X_0 + \int_0^t X_{s-}^* Q^0 ds + \mathcal{M}_t$

[1] The tuple $(\boldsymbol{Y}, \boldsymbol{Z}, \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}, \mathbb{Q})$ solves (3) if $\boldsymbol{Y} \in \mathcal{H}^2$, $\boldsymbol{Z} \in \mathcal{H}^2_X$, $\boldsymbol{\alpha} \in \mathbb{A}$, $\boldsymbol{\rho} \in M(\mathcal{R})$, $\mathbb{Q}$ is a probability measure on $(\Omega, \mathcal{F})$ and (3) is satisfied $\mathbb{P} - a.s.$ for all $t \in [0, T]$.

$\mathcal{H}^2$ càdlàg, real-valued, $\mathbb{F}$-adapted $\boldsymbol{Y}$: $\mathbb{E}[\int_0^T Y_t^2 dt] < +\infty$

$\mathcal{H}^2_X$ left cont., $\mathbb{R}^m$-valued, $\mathbb{F}$-adapted $\boldsymbol{Z}$: $\mathbb{E}[\int_0^T \|Z\|^2_{X_{t-}} dt] < +\infty$

$\|z\|^2_{X_{t-}} = z^* \psi_t z$, $z \in \mathbb{R}^m$

Characterizing mean-field Nash equilibria 3/3

Hypothesis A

- The transition rates are bounded and Lipschitz continuous in control and law
- The running cost is Lipschitz continuous in control and law
- The Hamiltonian admits a unique minimizer which is
  - feedback in (t, z, ρ)
  - measurable
  - Lipschitz continuous in z

Proposition

Assume that Hypothesis A holds true and $(\boldsymbol{\lambda}, \xi)$ fixed and admissible.

- *If the FBSDE admits a solution $(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{\alpha}, \boldsymbol{\rho}, \mathbb{Q})$, then $(\boldsymbol{\alpha}, \boldsymbol{\rho}) \in \mathcal{N}(\boldsymbol{\lambda}, \xi)$.*
- *If $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}) \in \mathcal{N}(\boldsymbol{\lambda}, \xi)$, then the FBSDE admits a solution $(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{\alpha}, \boldsymbol{\rho}, \mathbb{Q})$ such that $\boldsymbol{\alpha} = \hat{\boldsymbol{\alpha}}$, $d\mathbb{P} \otimes dt$-a.s., and $\rho_t = \hat{\rho}_t$, dt-a.e.*

Proof.

Along the lines of Carmona-Wang (2018). □

Stackelberg game for epidemics with contract factor control 1/2

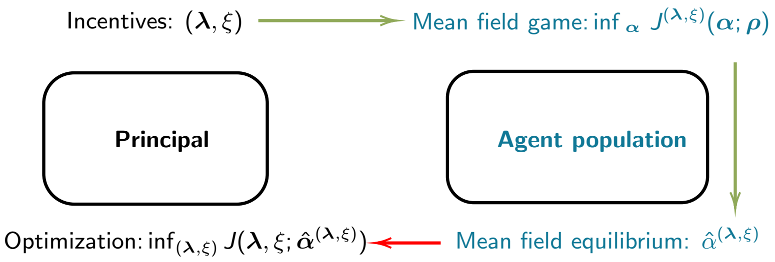

What is a Stackelberg game? Generically:

- 2 players: "leader" (principal) and "follower" (Mean field game)
- The leader moves first, then the follower moves
- The follower optimizes her objective function (finds the equilibrium) knowing the leaders move (the policy/incentive structure)
- The leader optimizes her objective function by anticipating the optimal (equilibrium) response from the follower

MFG for epidemics with contract factor control 2/2

Definition

A policy $(\boldsymbol{\lambda}, \xi)$ is admissible if $\boldsymbol{\lambda} \in \Lambda$ [2], ξ is $\mathcal{F}$-measurable, and $\mathcal{N}(\boldsymbol{\lambda}, \xi)$ is a singleton. We denote the set of admissible policies by $\mathcal{C}$.

[2] Λ: the set of measurable $\mathbb{R}_+^m$-valued functions with domain $[0, T]$

MFG for epidemics with contract factor control 2/2

Definition

A policy $(\boldsymbol{\lambda}, \xi)$ is admissible if $\boldsymbol{\lambda} \in \Lambda$ [2], ξ is $\mathcal{F}$-measurable, and $\mathcal{N}(\boldsymbol{\lambda}, \xi)$ is a singleton. We denote the set of admissible policies by $\mathcal{C}$.

The principal's cost for policy $(\boldsymbol{\lambda}, \xi) \in \mathcal{C}$ is

$$J(\boldsymbol{\lambda}, \xi) := \mathbb{E}^{\mathbb{Q}^{\mathcal{N}(\boldsymbol{\lambda}, \xi)}} \left[\int_0^T \left(c_0(t, \hat{\rho}_t^{\boldsymbol{\lambda}, \xi}(A, \cdot)) + f_0(t, \lambda_t) \right) dt + C_0(\hat{\rho}_T^{\boldsymbol{\lambda}, \xi}(A, \cdot)) + \xi \right]$$

[2] Λ: the set of measurable $\mathbb{R}_+^m$-valued functions with domain $[0, T]$

MFG for epidemics with contract factor control 2/2

Definition

A policy $(\boldsymbol{\lambda}, \xi)$ is admissible if $\boldsymbol{\lambda} \in \Lambda$ [2], ξ is $\mathcal{F}$-measurable, and $\mathcal{N}(\boldsymbol{\lambda}, \xi)$ is a singleton. We denote the set of admissible policies by $\mathcal{C}$.

The principal's cost for policy $(\boldsymbol{\lambda}, \xi) \in \mathcal{C}$ is

$$J(\boldsymbol{\lambda}, \xi) := \mathbb{E}^{\mathbb{Q}^{\mathcal{N}(\boldsymbol{\lambda}, \xi)}} \left[\int_0^T \Big(c_0(t, \hat{\rho}_t^{\boldsymbol{\lambda}, \xi}(A, \cdot)) + f_0(t, \lambda_t) \Big) \, dt + C_0(\hat{\rho}_T^{\boldsymbol{\lambda}, \xi}(A, \cdot)) + \xi \right]$$

If the population's equilibrium cost is too high, they **reject** the policy!

[2] Λ: the set of measurable $\mathbb{R}_+^m$-valued functions with domain $[0, T]$

MFG for epidemics with contract factor control 2/2

Definition

A policy $(\boldsymbol{\lambda}, \xi)$ is admissible if $\boldsymbol{\lambda} \in \Lambda$ [2], ξ is $\mathcal{F}$-measurable, and $\mathcal{N}(\boldsymbol{\lambda}, \xi)$ is a singleton. We denote the set of admissible policies by $\mathcal{C}$.

The principal's cost for policy $(\boldsymbol{\lambda}, \xi) \in \mathcal{C}$ is

$$J(\boldsymbol{\lambda}, \xi) := \mathbb{E}^{\mathbb{Q}^{\mathcal{N}(\boldsymbol{\lambda},\xi)}} \left[\int_0^T \Big(c_0(t, \hat{\rho}_t^{\boldsymbol{\lambda},\xi}(A, \cdot)) + f_0(t, \lambda_t) \Big) \, dt + C_0(\hat{\rho}_T^{\boldsymbol{\lambda},\xi}(A, \cdot)) + \xi \right]$$

If the population's equilibrium cost is too high, they **reject** the policy!

- Rejection whenever cost exceeds the **reservation threshold** $\kappa \in \mathbb{R}$
- The principal **disregards** policies that will be rejected
- The principal's **optimization problem** is

$$V(\kappa) := \inf_{(\boldsymbol{\lambda},\xi) \in \mathcal{C}} \; \inf_{\substack{(\boldsymbol{\alpha},\boldsymbol{\rho}) \in \mathcal{N}(\boldsymbol{\lambda},\xi) \\ J^{\boldsymbol{\lambda},\xi}(\boldsymbol{\alpha};\boldsymbol{\rho}) \leq \kappa}} J(\boldsymbol{\lambda}, \xi).$$

Holmström-Milgrom (1987), Sannikov (2008, 2013), Djehiche-Helgesson (2014), Cvitanić *et al* (2018), Carmona-Wang (2018), Elie *et al* (2019)

[2] Λ: the set of measurable $\mathbb{R}_+^m$-valued functions with domain $[0, T]$

Numerical approach to the Stackelberg game 1/9

How can we treat the Stackelberg game problem numerically?

Rewriting → **Discretization** → **Parameterization**

- Reposing the FBSDE as a control problem. "Sannikov's trick".
- Time-discretization and Monte Carlo-approximation.
- Parametrizing the optimization variables. Neural networks.

Numerical approach to the Stackelberg game 2/9

Rewriting → **Discretization** → **Parameterization**

Given $\boldsymbol{Z} \in \mathcal{H}^2_X$, $\boldsymbol{\lambda} \in \Lambda$, and real-valued $\mathcal{F}_0$-measurable Y_0, consider under $\mathbb{P}$:

$$\begin{cases} Y_t^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0} = Y_0 - \displaystyle\int_0^t \hat{H}(s, X_{s-}, Z_s, \hat{\rho}_s^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0})ds + \int_0^t Z_s^* \, d\mathcal{M}_s, \\ \qquad\quad \mathcal{E}_t = 1 + \displaystyle\int_0^t \mathcal{E}_{s-} X_{s-}^* \left(Q(s, \hat{\alpha}_s^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}, \hat{\rho}_s^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}) - Q^0 \right) \psi_s^+ \, d\mathcal{M}_s, \\ \hat{\rho}_t^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0} = \mathbb{Q}^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0} \circ \left(\hat{\alpha}_t^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}, X_t \right)^{-1}, \quad \dfrac{d\mathbb{Q}^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}}{d\mathbb{P}} = \mathcal{E}_T, \\ \hat{\alpha}_t^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0} = \hat{\alpha}(t, X_{t-}, Z_t, \hat{\rho}_t^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}). \end{cases}$$

Same equations as the FBSDE, except that the dynamic of $\boldsymbol{Y}$ is written in the **forward direction of time**.

Hypotesis B

- The function $U : \mathbb{R} \to \mathbb{R}$ is invertible.
- c_0, f_0 are measurable on $[0, T] \times \mathbb{R}^m$.

Consider the following optimal control problem

$$\widetilde{V}(\kappa) := \inf_{Y_0:\mathbb{E}[Y_0]\leq\kappa} \inf_{\substack{\boldsymbol{Z}\in\mathcal{H}^2_X \\ \boldsymbol{\lambda}\in\Lambda}} \mathbb{E}^{\mathbb{Q}^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}} \Bigg[\int_0^T \left(c_0\left(t, \hat{p}_t^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}\right) + f_0(t, \lambda_t) \right) dt + C_0\left(\hat{p}_T^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}\right) + U^{-1}\left(-Y_T^{\boldsymbol{Z},\boldsymbol{\lambda},Y_0}\right) \Bigg],$$

Proposition

If Hypothesis A and B then $\widetilde{V}(\kappa) = V(\kappa)$.

Proof.

Along the lines of Carmona-Wang (2018). □

- The backward equation has been "replaced" by an optimization problem.

Numerical approach to the Stackelberg game 4/9

Final polishing: express $\boldsymbol{Y}^{Z,\boldsymbol{\lambda},Y_0}$ with respect to $\mathcal{M}^{Z,\boldsymbol{\lambda},Y_o}$:

$$\begin{cases} Y_t^{Z,\boldsymbol{\lambda},Y_0} = Y_0 - \int_0^t f(s, X_{s-}, \hat{\alpha}_s^{Z,\boldsymbol{\lambda},Y_0}, \hat{\rho}_s^{Z,\boldsymbol{\lambda},Y_0}; \lambda_s) ds + \int_0^t Z_s^* d\mathcal{M}_s^{Z,\boldsymbol{\lambda},Y_0}, \\ \mathcal{E}_t = 1 + \int_0^t \mathcal{E}_{s-} X_{s-}^* \left(Q(s, \hat{\alpha}_s^{Z,\boldsymbol{\lambda},Y_0}, \hat{\rho}_s^{Z,\boldsymbol{\lambda},Y_0}) - Q^0 \right) \psi_s^+ d\mathcal{M}_s, \\ \hat{\rho}_t^{Z,\boldsymbol{\lambda},Y_0} = \mathbb{Q}^{Z,\boldsymbol{\lambda},Y_0} \circ \left(\hat{\alpha}_t^{Z,\boldsymbol{\lambda},Y_0}, X_t \right)^{-1}, \quad \frac{d\mathbb{Q}^{Z,\boldsymbol{\lambda},Y_0}}{d\mathbb{P}} = \mathcal{E}_T, \\ \hat{\alpha}_t^{Z,\boldsymbol{\lambda},Y_0} = \hat{\alpha}(t, X_{t-}, Z_t, \hat{\rho}_t^{Z,\boldsymbol{\lambda},Y_0}) \end{cases} \tag{4}$$

where the process $\boldsymbol{\mathcal{M}}^{Z,\boldsymbol{\lambda},Y_0}$ is defined by:

$$\mathcal{M}_t^{Z,\boldsymbol{\lambda},Y_0} = \mathcal{M}_t - \int_0^t X_{s-}^* \left(Q(s, \hat{\alpha}_s^{Z,\boldsymbol{\lambda},Y_0}, \hat{\rho}_s^{Z,\boldsymbol{\lambda},Y_0}) - Q^0 \right) ds,$$

is a $\mathbb{Q}^{Z,\boldsymbol{\lambda},Y_0}$-martingale. Furthermore, under $\mathbb{Q}^{Z,\boldsymbol{\lambda},Y_0}$,

$$X_t = X_0 + \int_0^t X_{s-}^* Q(s, \hat{\alpha}_s^{Z,\boldsymbol{\lambda},Y_0}, \hat{\rho}_s^{Z,\boldsymbol{\lambda},Y_0}) ds + \mathcal{M}_t^{Z,\boldsymbol{\lambda},Y_0}. \tag{5}$$

Numerical approach to the Stackelberg game 5/9

Rewriting → **Discretization** → **Parameterization**

Recall: the tuple $(\boldsymbol{Y}, \boldsymbol{Z}, \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}, \mathbb{Q})$ solves the FBSDE.

Hypothesis C

- $\hat{\alpha}$ depends only on the **state marginal** of the joint distribution: $\exists\ \check{\alpha} : [0, T] \times E \times \mathbb{R}^m \times \mathcal{P}(E) \to \mathbb{R}$ such that

$$\hat{\alpha}_t = \check{\alpha}(t, X_{t-}, Z_t, \hat{p}_t)$$

where $\hat{p}_t(\cdot) = \hat{\rho}_t(A, \cdot)$.

Hypothesis C is **weaker** than
assuming that $\hat{a}$ (the function) is independent of the first marginal of $\hat{\rho}$
(cf. Carmona-Wang (2018), Laurière-Tangpi (2019, 2020))

Numerical approach to the Stackelberg game 6/9

Input: Transition rate matrix function Q; number of particles N; time horizon T; initial distribution p_0; control functions λ, y_0, z
Output: Sampled trajectories for (4)–(5) (rewritten FBSDE)

Let $n = 0,\ t_0 = 0$; pick $X_0^i \sim p^0$ i.i.d and set $Y_0^i = y_0(X_0^i), i \in [\![N]\!]$
while $t_n \leq T$ **do**
  Set $Z_{t_n}^i = z(t_n, X_{t_n}^i), \alpha_{t_n}^i = \check{a}(t_n, X_{t_n}^i, Z_{t_n}^i, p_{t_n}), i \in [\![N]\!]$
  Let $\bar{\rho}_{t_n}^N = \frac{1}{N}\sum_{i=1}^N \delta_{(X_{t_n}^i, \alpha_{t_n}^i)}$ and $\bar{p}_{t_n}^N = \frac{1}{N}\sum_{i=1}^N \delta_{X_{t_n}^i}$
  Pick $(T^{i,e})_{e\in E, i\in [\![N]\!]}$ i.i.d. with exponential distribution of parameter 1
  Set the holding times: $\tau^{i,e} = T^{i,e}/Q_{X_{t_n}^i, e}(t_n, \alpha_{t_n}^i, \bar{\rho}_{t_n}^N),\ i \in [\![N]\!], e \in E$
  Let $e_\star^i =_{e\in E} \tau^{i,e}$ and $\tau_\star^i = \tau^{i, e_\star^i} = \min_{e\in E} \tau^{i,e}, i \in [\![N]\!]$
  Let $i_\star =_{i\in [\![N]\!]} \tau_\star^i$ be the first particle to jump
  Let $\Delta t = \tau_\star^{i_\star}$; set $X_{t_n+\Delta t}^{i_\star} = e_\star^{i_\star}$, and for every $i \neq i_\star$, set $X_{t_n+\Delta t}^i = X_{t_n}^i$
  Let $\Delta M_{t_n}^i = X_{t_n+\Delta t}^i - X_{t_n}^i - (X_{t_n}^i)^* Q(t_n, \alpha_{t_n}^i, \bar{\rho}_{t_n}^N)\Delta t,\ i \in [\![N]\!]$
  Let $Y_{t_n+\Delta t}^i = Y_{t_n}^i - f(t, X_{t_n}^i, \alpha_{t_n}^i, \bar{\rho}_{t_n}^N; \lambda(t_n))\Delta t + (Z_{t_n}^i)^* \Delta M_{t_n}^i,\ i \in [\![N]\!]$
  Set $n = n + 1$ and $t_n = t_{n-1} + \Delta t$
end while
Set $n_{tot} = n, t_{n_{tot}} = T, (X_{t_{n_{tot}}}^i, Y_{t_{n_{tot}}}^i, Z_{t_{n_{tot}}}^i) = (X_{t_{n_{tot}-1}}^i, Y_{t_{n_{tot}-1}}^i, Z_{t_{n_{tot}-1}}^i)$
return $(X_{t_n}^i, Y_{t_n}^i, Z_{t_n}^i)_{n=0,\dots,n_{tot}, i\in [\![N]\!]}$ and $(t_n)_{n=0,\dots,n_{tot}}$

Numerical approach to the Stackelberg game 7/9

Rewriting → **Discretization** → **Parameterization**

- Parameterize $(\boldsymbol{Z}, \boldsymbol{\lambda}, Y_0)$

$$z_{\theta_1} : [0, T] \times E \to \mathbb{R}^m, \quad \lambda_{\theta_2} : [0, T] \to \mathbb{R}^m_+, \quad y_{0,\theta_3} : E \to \mathbb{R}$$

- Feedforward fully connected neural networks
- The principal's cost for $(\theta_1, \theta_2, \theta_3)$:

$$\mathbb{J}^N(\theta) = \frac{1}{M} \sum_{j=1}^{M} \left[\sum_{n=0}^{n_{tot}-1} \left(c_0 \left(t_n, \bar{p}^{j,N,\theta}_{t_n} \right) + f_0(t_n, \lambda_{\theta_2}(t_n)) \right) (t_{n+1} - t_n) \right.$$
$$\left. + \, C_0 \left(\bar{p}^{j,N,\theta}_T \right) + \frac{1}{N} \sum_{i=1}^{N} U^{-1} \left(-Y^{j,i,\theta}_T \right) \right],$$

where for $j = 1, \dots, M$, $(\boldsymbol{Y}^{j,i,\theta})_{i \in [\![N]\!]}$ and $\bar{\boldsymbol{p}}^{j,N,\theta}$ are constructed using $(z, \lambda, y_0) = (z_{\theta_1}, \lambda_{\theta_1}, y_{0,\theta_3})$.

Numerical approach to the Stackelberg game 8/9

Final goal: minimize $\mathbb{J}^N$ over NN parameters $\theta = (\theta_1, \theta_2, \theta_3)$.

Minimization by **Adaptive Moment Estimation algorithm**:

- second algorithm of Carmona-Laurière (2019)
- adapted to
  - the finite state case
  - the Stackelberg setting

For a sample $S = (X^i_{t_n}, Y^i_{t_n}, Z^i_{t_n})_{n=0,\dots,n_{tot}, i\in[\![N]\!]}$:

$$\mathbb{J}^N_S(\theta) = \sum_{n=0}^{n_{tot}-1} \left(c_0\left(t_n, \bar{p}^{N,\theta}_{t_n}\right) + f_0(t_n, \lambda_{\theta_2}(t_n)) \right)(t_{n+1} - t_n) \\ + C_0\left(\bar{p}^{N,\theta}_T\right) + \frac{1}{N}\sum_{i=1}^{N} U^{-1}\left(-Y^{i,\theta}_T\right) \tag{6}$$

where $\overline{p}^N_{t_n} = \frac{1}{N}\sum_{i=1}^N \delta_{X^i_{t_n}}$.

Numerical approach to the Stackelberg game 9/9

Input: Initial parameter θ_0; number of iterations K; sequence $(\beta_k)_{k=0,\dots,K-1}$ of learning rates; transition rate matrix function Q; number of particles N; time horizon T; initial distribution p_0
Output: Approximation of θ^* minimizing $\mathbb{J}^N$

for $k = 0, 1, 2, \dots, K-1$ **do**
  Sample $S = (X^i_{t_n}, Y^i_{t_n}, Z^i_{t_n})^{i \in [\![N]\!]}_{n=0,\dots,n_{tot}}$ and $(t_n)_{n=0,\dots,n_{tot}}$ with controls $(z, \lambda, y_0) = (z_{\theta_{k,0}}, \lambda_{\theta_{k,1}}, y_{0,\theta_{k,2}})$ and parameters: Q, N, T, p_0
  Compute the gradient $\nabla \mathbb{J}^N_S(\theta_k)$ of $\mathbb{J}^N_S(\theta_k)$ defined by (6)
  Set $\theta_{k+1} = \theta_k - \beta_k \nabla \mathbb{J}^N_S(\theta_k)$
end for
return θ_K

Example: SIR MFG and an inactive principal 1/6

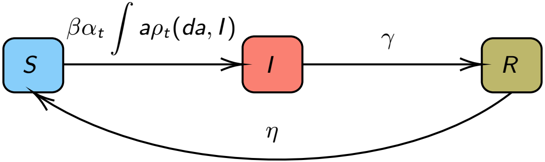

$$f(t,x,\alpha,\rho;\lambda) = \frac{c_\lambda}{2}\left(\lambda^{(S)}-\alpha\right)^2 \mathbb{1}_S(x) + \left(\frac{1}{2}\left(\lambda^{(I)}-\alpha\right)^2 + c_I\right)\mathbb{1}_I(x) + \frac{1}{2}\left(\lambda^{(R)}-\alpha\right)^2 \mathbb{1}_R(x), \tag{7}$$

- Deviation from recommended contact factor λ
- Infection cost

Example: SIR MFG and an inactive principal 2/6

Hypothesis D (to have semi-explicit solutions!)

- There exists a unique solution $(\hat{\boldsymbol{Y}}, \hat{\boldsymbol{Z}}, \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}, \hat{\mathbb{Q}})$ to the FBSDE.
- Evaluated at the equilibrium, $\hat{a}$, f, and Q are functions of the state-marginal law only: $\bar{a}, \bar{f}, \bar{Q}$.
- The function $\bar{a}$ is Lischitz continuous in z and p (the state marginal).

Example: SIR MFG and an inactive principal 2/6

Hypothesis D (to have semi-explicit solutions!)

- There exists a unique solution $(\hat{\boldsymbol{Y}}, \hat{\boldsymbol{Z}}, \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}, \hat{\mathbb{Q}})$ to the FBSDE.
- Evaluated at the equilibrium, $\hat{a}$, f, and Q are functions of the state-marginal law only: $\bar{a}, \bar{f}, \bar{Q}$.
- The function $\bar{a}$ is Lischitz continuous in z and p (the state marginal).

Definition

Let $(\boldsymbol{\alpha}, \boldsymbol{p}) \in \mathbb{A} \times M(\mathcal{P}(E))$ and denote by $\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{p}} \in \mathcal{P}(\Omega)$ the measure such that the coordinate process X_t has transition rate matrix $\bar{Q}(t, \alpha_t, p_t)$ under $\mathbb{Q}^{\boldsymbol{\alpha},\boldsymbol{p}}$. Assume that $(\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{p}}) \in \mathbb{A} \times M(\mathcal{P}(E))$ satisfies

(i) $\bar{\boldsymbol{\alpha}} = \arg\inf_{\alpha \in \mathbb{A}} \mathbb{E}^{\mathbb{Q}^{\boldsymbol{\alpha},\bar{\boldsymbol{p}}}} \left[\int_0^T \bar{f}(t, X_t, \alpha_t, \bar{p}_t) dt - U(\xi) \right]$,

(ii) $\forall t \in [0, T], i \in \{1, \ldots, m\}: \; \bar{p}_t(i) = \mathbb{Q}^{\bar{\boldsymbol{\alpha}},\bar{\boldsymbol{p}}}(X_t = e_i)$.

Then $(\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{p}})$ is called a **non-extended mean field Nash equilibrium**.

Proposition

Assume Hypothesis A–D to be true. Denote the tuple of Hypothesis D by $(\hat{\boldsymbol{Y}}, \hat{\boldsymbol{Z}}, \hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}}, \mathbb{Q})$. The pair $(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\rho}})$ is a mean-field Nash equilibrium. Let $\hat{p}_t$ be the E-marginal of $\hat{\rho}_t$ and let $(\bar{\boldsymbol{\alpha}}, \bar{\boldsymbol{p}})$ be a non-extended mean field Nash equilibrium. Then $\hat{p}_t = \bar{p}_t$ for dt-a.e. $t \in [0, T]$ and $\hat{\alpha}_t = \bar{\alpha}_t$ $d\mathbb{P} \otimes dt$-a.e..

Example: SIR MFG and an inactive principal 3/6

The regulator declares a **fixed policy** $(\boldsymbol{\lambda}, \xi)$

Test case	**Contact factor**	ξ	$\lambda_t^{(S)}$	$\lambda_t^{(I)}$	$\lambda_t^{(R)}$
Free spread	Constant	0	1	1	1
No lockdown	MF Nash eq.	0	1	1	1
Late lockdown	MF Nash eq.	0	$1 - 0.3\mathbb{1}_{t>40}$	$0.9 - 0.3\mathbb{1}_{t>40}$	1
Early lockdown	MF Nash eq.	0	$1 - 0.3\mathbb{1}_{t\leq 10}$	$0.9 - 0.3\mathbb{1}_{t\leq 10}$	1

Parameter choice

Parameter	T	p^0	c_λ	c_I	β	γ	η
Value in tests	50	$(0.9, 0.1, 0)$	10	1	0.25	0.1	0

Example: SIR MFG and an inactive principal 4/6

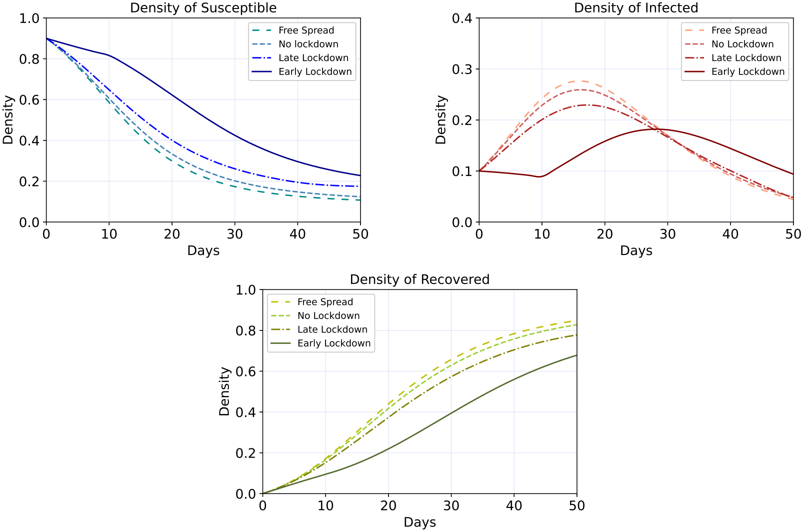

Figure: Semi-explicit (ODE) solution in the four test cases

Example: SIR MFG and an inactive principal 5/6

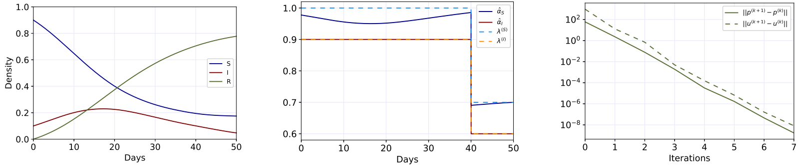

Figure: Late lockdown, ODE solution. Evolution of the population state distribution (left), evolution of the controls (middle), convergence of the solver (right).

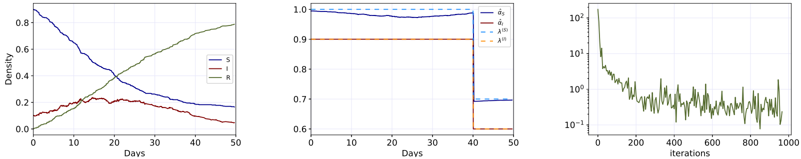

Figure: Late lockdown, numerical solution. Evolution of the population state distribution (left), evolution of the controls (middle), convergence of the loss value (right).

Example: SIR MFG and an inactive principal 6/6

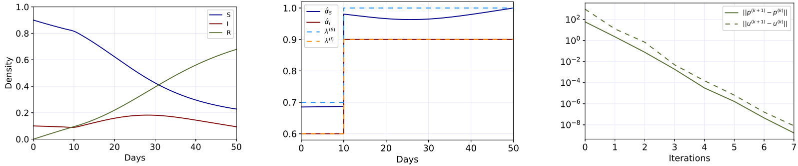

Figure: Early lockdown, ODE solution. Evolution of the population state distribution (left), evolution of the controls (middle), convergence of the solver (right).

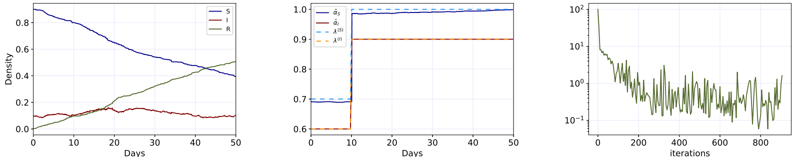

Figure: Early lockdown, numerical solution. Evolution of the population state distribution (left), evolution of the controls (middle), convergence of the loss value (right).

Example: SIR Stackelberg game 1/2

We now include the **regulator's optimization** to the previous example, making the problem a Stackelberg game.

More specifically, we set $C_0(p) = 0$ and

$$c_0(t, p) = c_{\mathrm{Inf}}\, p(I)^2, \quad f_0(t, \lambda) = \sum_{i \in \{S,I,R\}} \frac{\bar{\beta}^{(i)}}{2} \left(\lambda^{(i)} - \bar{\lambda}^{(i)}\right)^2 \tag{8}$$

for constant $\bar{\lambda}, \bar{\beta} \in \mathbb{R}_+^m$ and $c_{\mathrm{Inf}} > 0$.

- Deviation from some incentive levels $\bar{\lambda}$
- Infection cost

For this case we can derive a **semi-explicit solution**.

T	p^0	c_λ	c_I	c_{Inf}	$\bar{\beta}$	$\bar{\lambda}$	β	γ	η	κ
30	$(0.9, 0.1, 0)$	10	0.5	1	$(0.2, 1, 0)$	$(1, 0.7, 0)$	0.25	0.1	0	0

Example: SIR Stackelberg game 2/2

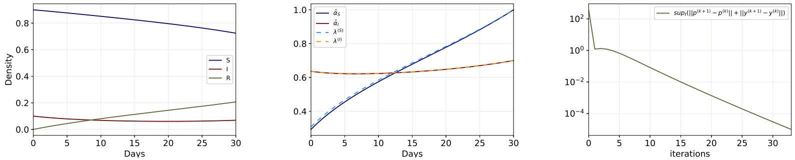

Figure: SIR Stackelberg game, ODE solution. Evolution of the population state distribution (left), evolution of the controls (middle), convergence of the solver (right).

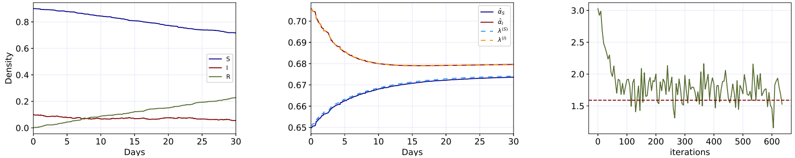

Figure: SIR Stackelberg, numerical solution. Evolution of the population state distribution (left), evolution of the controls (middle), convergence of the loss value (right).

Conclusions

A Stackelberg game to model decision making in an epidemic.

- The model incorporates at the same time a non-cooperative population and a regulator such as a government
- Evolution of the system described from the point of view of a typical (infinitesimal) agent
- Numerical method based on neural network approximation and Monte Carlo simulations to compute the optimal policy

What lies ahead?

- Further work on the FBSDE system to justify assumptions about its solution
- Generalizing beyond the SIR model is crucial for applications to epidemiological models. Multiple populations, pharmaceutical interventions, testing, etc.
- Other ways of modeling incentives

Thank you!