Optimal Vaccination Schedule

A Mean Field Game Approach

Josu Doncel[1] Nicolas Gast[2] Bruno Gaujal[3]

Bordeaux, June 28, 2019

[1]Univ. of Basque Country.
[2]Univ Grenoble Alpes and Inria.
[3]Univ Grenoble Alpes and Inria.

SIR Epidemic Model

A large number N of agents (nodes, persons, players) subject to interactions (they meet, communicate, ...).

Each agent has 3 possible states: Susceptible, Infected, Recovered (S, I, R).

When an agent in state S meets an agent in I, it gets infected.
An agent in state I will eventually recover and go to state R.
An agent in state R stays in R forever.

SIR Epidemic Model

Simplest epidemic model. Naive but with a good predictive power for human epidemies.

Introduced in 1927 by Kermack–McKendrick.
"*Because of their seminal importance to the field of theoretical epidemiology, these articles were republished in the Bulletin of Mathematical Biology in 1991*." [From wikipedia].

Has been studied ever since, 100s of papers in mathematics, computer science, health studies, bio-informatics.

SIRV Dynamics

- A player encounters other players with rate γ (activity of the player). If the first player is Susceptible and the second is Infected, the first one becomes Infected.
- An Infected player Recovers at rate ρ.
- A Susceptible player can decide to get vaccinated with rate $\pi(t) \in [0, M]$.
- Once a player is vaccinated or recovered, its state (R) does not change.

Let $(m_S(t), m_I(t), m_R(t)$ the proportion of the players in states S, I, R. The Markovian evolution of one player is

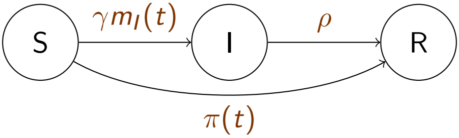

Cost Functions and Objectives

The cost of being infected is c_I per time unit.

The vaccination cost is linear in the rate π of the vaccination chosen by the player: $c_V \cdot \pi$.

Problem to be solved : *under full information (state of all players at time t and their vaccination schedule is known to all), each player wants to choose a vaccination schedule (strategy) that minimizes its cost up to a time horizon T.*

Cost Functions and Objectives

The cost of being infected is c_I per time unit.

The vaccination cost is linear in the rate π of the vaccination chosen by the player: $c_V \cdot \pi$.

Problem to be solved : *under full information (state of all players at time t and their vaccination schedule is known to all), each player wants to choose a vaccination schedule (strategy) that minimizes its cost up to a time horizon T.*

Not a well-posed problem: the optimal schedule of a player depends on the schedule of any other player, who in turn is trying to optimize its vaccination schedule that depends on the first player's schedule.

Nash Equilibrium and Social Optimal

Definition (Nash Equilibrium (NE))

A Nash Equilibrium is a vaccination schedule π_{NE} such that if all the players use π_{NE}, then any player's optimal strategy is to use π_{NE}.

Definition (Social Optimal)

A social optimal is a vaccination schedule π_{SO} that minimizes the sum of the costs of all the players.

NE always exist in SIRV (Kakutani fixed point theorem).
SO always exist in SIRV (compacity of the strategy space for weak topology).

Unfortunately both are very hard to compute when N is large (combinatorial explosion of the state space).

Mean Field Model (Fluid Model)

Since the players are all the same, the state of the system is given by the population distribution, $(m_S(t), m_I(t), m_R(t))$.

Mean Field Model (Fluid Model)

Since the players are all the same, the state of the system is given by the population distribution, $(m_S(t), m_I(t), m_R(t))$.
When $N \to \infty$, $(m_S(t), m_I(t), m_R(t))$ behaves as a fluid, whose evolution follows the Kolmogorov equations of the individual Markov chain. Under vaccination strategy π,

$$\begin{cases} \dot{m}_S(t) = -\gamma m_S(t) m_I(t) - \pi(t) m_S(t) \\ \dot{m}_I(t) = \gamma m_S(t) m_I(t) - \rho m_I(t) \\ \dot{m}_R(t) = \rho m_I(t) + \pi(t) m_S(t). \end{cases}$$

Mean Field Model (Fluid Model)

Since the players are all the same, the state of the system is given by the population distribution, $(m_S(t), m_I(t), m_R(t))$.
When $N \to \infty$, $(m_S(t), m_I(t), m_R(t))$ behaves as a fluid, whose evolution follows the Kolmogorov equations of the individual Markov chain. Under vaccination strategy π,

$$\begin{cases} \dot{m}_S(t) = -\gamma m_S(t) m_I(t) - \pi(t) m_S(t) \\ \dot{m}_I(t) = \gamma m_S(t) m_I(t) - \rho m_I(t) \\ \dot{m}_R(t) = \rho m_I(t) + \pi(t) m_S(t). \end{cases}$$

Some technicalities here...$\pi(t)$ may not be continuous...(Carathéodory Existence Theorem).

When $\pi = 0$, this is classical SIR dynamics Kermack–McKendrick (1927). Analytical solution derived recently by Harko, Lobo and Mak (2014).

Strategy for One Player

Player 0 using strategy π^0 while the population uses π. The state probabilities (p_S^0, p_I^0, p_R^0) of Player 0 has an evolution given by its local Kolmogorov equation:

$$\begin{cases} \dot{p}_S^0(t) = -\gamma p_S^0(t) m_I(t) - \pi^0(t) p_S^0(t) \\ \dot{p}_I^0(t) = \gamma p_S^0(t) m_I(t) - \rho p_I^0(t) \\ \dot{p}_R^0(t) = \rho p_I^0(t) + \pi(t) p_S^0(t). \end{cases}$$

Using the foregoing notations, the expected individual cost of Player 0 is:

$$W(\pi^0, \pi) = \int_0^T \left(c_V \pi^0(t) p_S^0(t) + c_I p_I^0(t) \right) dt,$$

where c_V is the vaccination cost and c_I is the unit time cost of being infected.

Best Response Equation of One Player

The best response of Player 0 to a population using strategy π is a strategy π^0_* that minimizes its cost.

$W_X(t)$: optimal total cost from t to T of Player 0 when in state X at time t.
This defines a MDP whose Hamilton-Jacobi-Bellman equation is:

$$W_S(T) = W_I(T) = 0.$$

$$\begin{aligned} -\dot{W}_S(t) &= \inf_{\pi^0(t)} \left[\pi^0(t)\left(c_V - W_S(t)\right) + \gamma m_I(t)(W_I(t) - W_S(t))\right] \\ -\dot{W}_I(t) &= c_I - \rho W_I(t). \end{aligned}$$

$$\pi^0_*(t) = \arg\min_{\pi^0(t)} \left[\pi^0(t)\left(c_V - W_S(t)\right) + \gamma m_I(t)(W_I(t) - W_S(t))\right].$$

Mean Field Equilibria

Let π be the strategy used by the whole population.

Let $BR(\pi)$ be the best response (π^0_*) of Player 0 to π.

Definition (Mean Field Equilibirum)

If $\pi = BR(\pi)$ then π is a mean field equilibrum.

Mean Field Game theory developped initially by P.L. Lions (2007) in a more general framework.
Has had a large success in crowd movements, routing in telecommunication networks, stock markets... even sailing competition (MFG Labs)

Mean Field Equilibria for SIR

Lemma

For any population strategy π, there exists a best-response π_*^0 that is a threshold strategy: There exists a critical time t_c^0 s.t.

$$\pi_*^0(t) = M \quad t < t_c^0,$$
$$\pi_*^0(t) = 0 \quad t > t_c^0.$$

$$\pi_*^0(t) \;=\; \arg\min_{\pi^0(t)} \left[\pi^0(t)\left(c_V - W_S(t)\right) + \gamma m_I(t)(W_I(t) - W_S(t))\right].$$

Mean Field Equilibria for SIR

Lemma

For any population strategy π, there exists a best-response π^0_* that is a threshold strategy: There exists a critical time t^0_c s.t.

$$\pi^0_*(t) = M \quad t < t^0_c,$$
$$\pi^0_*(t) = 0 \quad t > t^0_c.$$

$$\pi^0_*(t) \quad = \quad \arg\min_{\pi^0(t)} \left[\pi^0(t)\left(c_V - W_S(t)\right) + \gamma m_I(t)(W_I(t) - W_S(t))\right].$$

Theorem

SIRV has a unique mean-field equilibrium, pure, with threshold t_c^{MFE}.

Same as Francis, 2004 (does not use the notion of MFE).

Social Optimal Strategy

We denote by $C(\pi)$ the total cost incurred by the population under strategy π, i.e.,

$$C(\pi) = \int_0^T (c_I m_I(t) + c_V \pi(t) m_S(t)) \, dt.$$

The global optimum of the problem is the population strategy that minimizes the total cost:

$$\pi^{opt} \in \arg\min_{\pi} C(\pi).$$

Using the Pontryagin maximum principle,

Proposition

The strategy that minimizes the total cost is a threshold strategy,

Social Optimal Strategy

We denote by $C(\pi)$ the total cost incurred by the population under strategy π, i.e.,

$$C(\pi) = \int_0^T (c_I m_I(t) + c_V \pi(t) m_S(t)) \, dt.$$

The global optimum of the problem is the population strategy that minimizes the total cost:

$$\pi^{opt} \in \arg\min_{\pi} C(\pi).$$

Using the Pontryagin maximum principle,

Proposition

The strategy that minimizes the total cost is a threshold strategy, with a larger threshold $t_c^{opt} \geq t_c^{MFG}$.

Numerical Comparisons

Using $\rho = 36.5$, $\gamma = 73$, $\tau = 10$, $T = 0.3$, $c_I = 36.5$ and $c_V = 0.5$ (typical for human epidemics).

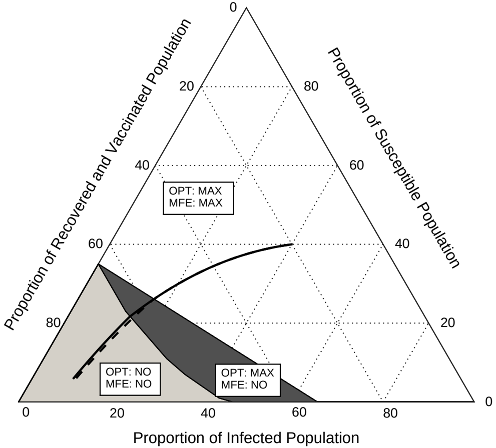

Population dynamics under the MFE (dashed) and social optimal (solid).

Numerical Comparisons (Zoom)

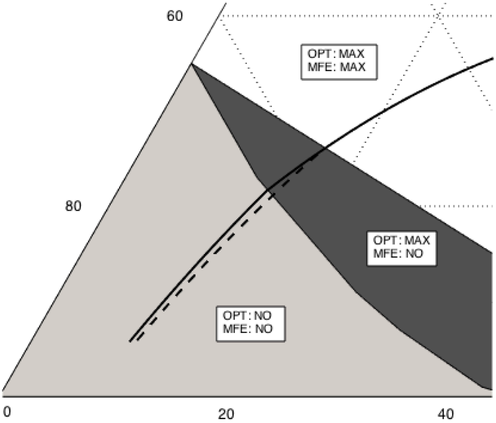

Mechanism Design

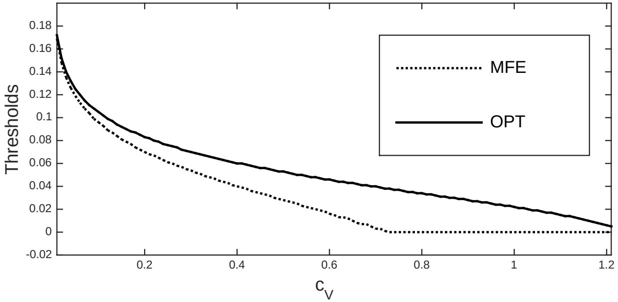

Thresholds of the MFE and of the global optimum when $c_V \in [0.01, 1.21]$.

If vaccination is let to individuals, then it should be subsidized to get a social optimal. Subsidizing by h (horizontal distance), both thresholds coincide:

$$t_c^{MFE}(c_V - h) = t_c^{opt}(c_V).$$

Comparisons with N Player Game

Optimal strategy for the N player game can be computed using an MDP approach (up to $N = 50$ using symmetry).
The cost of the NE converges to the cost of the MFE in $1/N$. (best known bound $1/\sqrt{N}$).

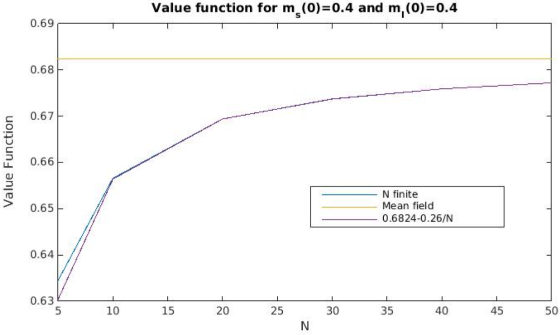

Figure: Cost of the Nash equilibirum with N players as N grows, and best fit of the form $a + b/N$.

Comparisons with N Player Game (II)

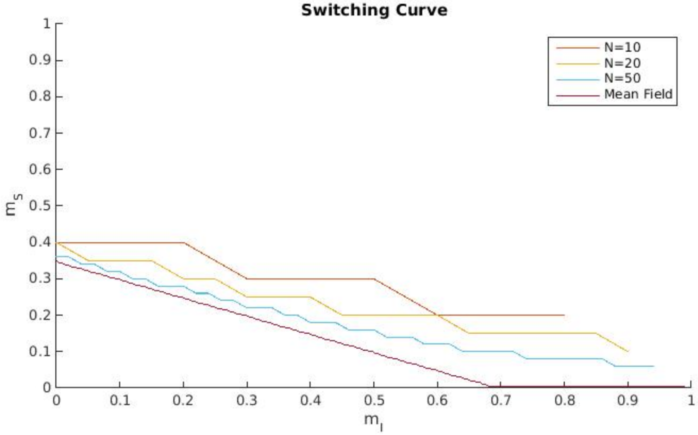

Figure: Convergence of the Nash switching curve to to the MFE switchwing curve.

What is missing?

- Proof of convergence rate in $1/N$.
- Partial information (local, signals,...).
- Neighboring graph.

The End.