

A Mean-Field Game Analysis of SIR Dynamics with Vaccination

Josu Doncel, Nicolas Gast, Bruno Gaujal

► **To cite this version:**

Josu Doncel, Nicolas Gast, Bruno Gaujal. A Mean-Field Game Analysis of SIR Dynamics with Vaccination. 2017. <hal-01496885>

HAL Id: hal-01496885

<https://hal.inria.fr/hal-01496885>

Submitted on 28 Mar 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Mean-Field Game Analysis of SIR Dynamics with Vaccination.

Josu Doncel^{a,b}, Nicolas Gast^{a,b}, Bruno Gaujal^{a,b}

^aUniv. Grenoble Alpes, F-38000 Grenoble, France

^bInria

Abstract

In this paper, we analyze a mean-field game model of SIR dynamics (Susceptible, Infected, Recovered) where players can vaccinate. This game admits a unique mean-field equilibrium: The equilibrium strategy of each player is to vaccinate until the proportion of susceptible players drops below some threshold and stop vaccinating thereafter. We also show that the vaccination strategy minimizing the total cost for the population is a different threshold strategy. This implies that, to encourage optimal vaccination behaviors, vaccination should be subsidized.

1. Introduction

Mean-field game theory studies the strategic decision-making of an infinite number of rational agents (or players) in interaction [10, 8, 9]. Two important assumptions are made in this type of models: (i) that the agents are indistinguishable, which implies that the strategy of an individual player only depends on the number of agents in each state (this also called the mean field assumption) and (ii) the focus is on symmetric strategies (all players in the same state adopt the same strategy).

Under these assumptions, mean-field games can be seen as approximations of finite symmetric games. They have attracted significant attention from the research community in the last decade, see for example [4, 1, 2]. The interest of mean-field games comes from the fact that they provide a simplification of the computation of the solution of the games, compared with games with a finite, but large, number of players. This paper is another illustration of this phenomenon. As all players have a finite state space, the equilibrium can be characterized by an ODE, the Bellman optimality equation. We show that this Bellman optimality equation has a single solution which implies the unicity and purity of the equilibrium.

More precisely, we explore one example of a mean-field game model of the spread of an epidemic where each player has 3 possible states. The dynamics of the population is given by the classical SIR (Susceptible-Infected-Recovered) model to which we add a control action for each player, namely vaccination. The susceptible population can get vaccinated, which has a one shot vaccination cost. Being infected incurs an infection cost per time unit.

In the mean-field game under study each individual player can choose how she gets vaccinated. The goal of an individual player is to minimize her expected cost. Potential applications include disease spreading in a large population, anti-virus usage in computer systems or on a dual vision, advertising systems, where the goal is to propagate information.

One can observe that our model is a particular case of the mean-field games introduced in [3], where the dynamics of one player depends explicitly on the population distribution. In [3], the authors show that there always exists a mean-field equilibrium in such games as soon as the dynamics and the cost are continuous functions. The existence of a solution of our mean-field game follows directly their result. Here, however, we go beyond this existence result and characterize the mean-field equilibrium. In [7] the authors develop an approximation of the same epidemic model by making independence assumptions, and they characterize the solution of the corresponding mean-field game. In this paper, we show that the mean-field game corresponding to this model is tractable without using independence approximations. We prove that there is a unique symmetric mean-field equilibrium that is a pure strategy consisting of vaccinating at maximum rate until a given threshold is reached and not to vaccinate thereafter.

We also study the efficiency of the mean-field equilibrium by comparing it with the centralized optimal control problem in this epidemic model, where the goal is to find a vaccination strategy for the whole population so as to minimize the total cost of the system (social cost). The solution of this problem is also of threshold type (it also consists of vaccinating at maximum rate until a given threshold is reached and to stop vaccinating after this). We compare the mean-field equilibrium and the global optimum and show that the threshold at which the global optimum stops vaccinating is always larger than that of the mean-field equilibrium. This suggests the following mechanism design: subsidizing individual vaccination can reduce the total cost of the mean-field equilibrium down to the optimal social cost.

The rest of the article is organized as follows. In Section 2, we describe the model and we define the notion of mean-field equilibrium. The main results of the mean-field game are given in Section 3. The centralized control problem is analyzed in Section 4 as well as the comparison with the mean-field equilibrium. The numerical analysis of both strategies as well as the numerical computations of subsidies are presented in Section 6.

^{*}This work is supported by the European FP7 Quanticol Project.

2. Model Description

We consider a population of homogeneous players that evolve in continuous time from 0 to a finite time horizon T . The players can be in one of the following states: susceptible (S), infected (I), recovered or vaccinated (R). We denote by $m_S(t)$, $m_I(t)$ and $m_R(t)$ the proportion of the population that is, respectively, susceptible, infected, recovered or vaccinated at time t .

The dynamics of one player is a Markov process in continuous time that can be described as follows. A player encounters other players with rate γ . If the initial player is susceptible and the encountered one is infected, the first player becomes infected. An infected player recovers at rate ρ . We also consider that a susceptible player can choose to get vaccinated with strategy $\pi(t)$, a measurable function from $[0, T]$ to $[0, \tau]$, where $\tau < \infty$ is the maximal vaccination rate. Once a player is vaccinated or has recovered, her state does not change. The Markovian behavior of a player is represented in Figure 1. Note that this is one of the simplest possible controlled Markov chain.




Figure 1: The dynamics of a player in the epidemic model. An player has three possible states: S (susceptible), I (infected), R (recovered or vaccinated).

The cost incurred by one player is assumed to be of the following form.

- Vaccination cost: The vaccination cost of a susceptible player is a linear function of her vaccination rate: $c_V\pi(t)$, where $c_V > 0$.
- Infection cost: Whenever a player is infected, it suffers a cost of c_I ($c_I > 0$) per time unit.

We are interested in the analysis of this epidemic model at the mean field limit. When the number of players goes to infinity, the dynamics of the population where all players use strategy π is given by a system of differential equations that is typical of population dynamics studied in [5]. Here, this system of differential equations is

$$\begin{cases} \dot{m}_S(t) = -\gamma m_S(t)m_I(t) - \pi(t)m_S(t) \\ \dot{m}_I(t) = \gamma m_S(t)m_I(t) - \rho m_I(t) \\ \dot{m}_R(t) = \rho m_I(t) + \pi(t)m_S(t). \end{cases} \quad (1)$$

2.1. Mean-Field Game

We focus on a particular player, that we call Player 0. Player 0 chooses her vaccination strategy π^0 , where $\pi^0(t) \in [0, \tau]$ for all t . The probability that Player 0 is in a given state depends not only on π^0 , but also on $m(t)$, the population distribution. The dynamics of the global population is still driven by Equations (1), because Player 0 alone does not modify the population distribution (m_S, m_I, m_R, m_V).

Let $x_i^{\pi^0, \pi}(t)$ be the probability that Player 0 is in state i at time t , where $i \in \{S, I, R\}$. The quantities $x_i^{\pi^0, \pi}(t)$ satisfy the following system of differential equations:

$$\begin{cases} \dot{x}_S^{\pi^0, \pi}(t) = -\gamma x_S^{\pi^0, \pi}(t)m_I(t) - \pi(t)x_S^{\pi^0, \pi}(t) \\ \dot{x}_I^{\pi^0, \pi}(t) = \gamma x_S^{\pi^0, \pi}(t)m_I(t) - \rho x_I^{\pi^0, \pi}(t) \\ \dot{x}_R^{\pi^0, \pi}(t) = \rho x_I^{\pi^0, \pi}(t) + \pi(t)x_S^{\pi^0, \pi}(t). \end{cases}$$

The expected individual cost of Player 0 is:

$$V^{\pi^0, \pi} = \int_0^T [c_V\pi^0(u)x_S^{\pi^0, \pi}(u) + c_I x_I^{\pi^0, \pi}(u)] du.$$

We call the *best response to π* and denote by $BR(\pi)$ the set of vaccination strategies that minimize the expected cost of Player 0 for a given population strategy π , that is,

$$BR(\pi) \in \arg \min_{\pi^0} V^{\pi^0, \pi}.$$

The set $BR(\pi)$ is not empty by compactity of the strategy space¹. Notice that the best response correspondence in the mean field limit is simpler to define than the best response of a stochastic symmetric game with a finite number of players. Indeed, with a finite set of players, the strategy of Player 0 modifies the state of Player 0 but also the global state of the system. In turn, under their policy π , all other players react to the choices of Player 0 so that their state will also depend on π^0 , making the computation of the best response correspondence quite complex.

Now, by definition, a mean-field equilibrium is a vaccination strategy π^{MFE} such that, when the population applies strategy π^{MFE} , the best response of Player 0 is the strategy π^{MFE} . In other words, a mean-field equilibrium is a fixed point for the best-response function.

Definition 1 (Symmetric Mean-Field Equilibrium). *The vaccination policy π^{MFE} is a symmetric mean-field equilibrium if and only if*

$$\pi^{MFE} \in BR(\pi^{MFE}).$$

3. Mean-Field Equilibrium Characterization

In this section, we analyze the mean-field equilibrium described in Section 2.1. Before going further, we observe that the rate at which susceptible population becomes infected is linear in the proportion of infected population and the rest of the rates and costs do not depend on the population distribution. Thus, this model is a particular case of the mean-field games whose equilibrium existence is proven in [3]. As a result, to prove the existence of a mean-field equilibrium, one could invoke the general result of [3]. In this section, we use a more direct method to establish the existence of an equilibrium. This

¹The set of strategies is the set of bounded measurable functions on $[0, T]$ endowed with the weak topology. It is a compact set.

methods further implies that the equilibrium is unique and pure, and shows that the equilibrium is of threshold type.

From the discussion in Section 2.1, the minimization problem that consists in computing the best-response of Player 0 to the population can be seen as a continuous time Markov decision process with finite horizon T . Let us denote by $J_S(t)$ (resp. $J_I(t)$) the optimal cost to go, starting in state S (resp. I), from time t . The optimal costs and the optimal strategy π^{0*} of Player 0 satisfy the continuous time Bellman optimality equation (see [11]).

$$J_S(T) = J_I(T) = 0 \quad (2)$$

$$-\dot{J}_S(t) = \inf_{\pi^0} \left[\pi^0(t) (c_V - J_S(t)) + \gamma m_I(t) (J_I(t) - J_S(t)) \right] \quad (3)$$

$$-\dot{J}_I(t) = c_I - \rho J_I(t). \quad (4)$$

$$\begin{aligned} \pi^{0*}(t) &= \arg \min_{\pi^0} \left[\pi^0(t) (c_V - J_S(t)) + \gamma m_I(t) (J_I(t) - J_S(t)) \right] \\ &= \arg \min_{\pi^0} \left[\pi^0(t) (c_V - J_S(t)) \right] \end{aligned} \quad (5)$$

Proposition 1. *For any population strategy π , the best-response π^{0*} is a threshold strategy (up to a set of measure 0).*

Proof. First, let us remark that if the strategy of Player 0, $\pi^{0*}(t)$ is modified over a set of Lebesgue measure 0, then the costs and the states are not modified. This means that strategies are only defined up to a set of measure 0.

By (5), the best response policy $\pi^{0*}(t)$ is 0 if $J_S(t) < c_V$ and τ if $J_S(t) > c_V$. Therefore, to show the result, it suffices to show that $J_S(t)$ crosses c_V at most once.

Let t_1 be the first time when J_S gets below c_V . The solution of Equation (4) is

$$J_I(t) = \frac{c_I}{\rho} (1 - e^{\rho(t-T)}).$$

The cost $J_I(t)$ is decreasing from $\frac{c_I}{\rho}(1 - e^{-\rho T})$ to 0. As for J_S , as long as $J_S(t) < c_V$, then $\pi^{0*}(t) = 0$. At time T , $J_S(T) = 0 = J_I(T)$ so that $\dot{J}_S(T) = 0$. By continuity of the costs, this implies that $J_S(t) \leq J_I(t)$ for all $t_1 \leq t \leq T$ (also note that when $J_S(t) > c_V$, its derivative becomes smaller than when $J_S(t) < c_V$, so that $J_S(t) \leq J_I(t)$ for all $0 \leq t \leq T$).

This implies that $J_S(t)$ is decreasing between t_1 and T . Therefore, $J_S(t)$ can cross level c_V at most once (at time t_1). Otherwise, let us set $t_1 = 0$.

This implies that π^{0*} is a threshold policy:

$$\pi^{0*} = \begin{cases} \tau & \text{if } t < t_1 \\ 0 & \text{if } t \geq t_1. \end{cases} \quad (6)$$

□

Notice that the best response is not well-defined at the threshold point. This is in accordance with the fact that policies are only defined up to a set of measure 0.

Since the best response to any policy is threshold, we will only consider threshold policies in the following. Let us denote by $t_{BR(\theta)}$ the threshold of a best-response strategy for Player 0 to a population strategy π_θ , with threshold θ .

Lemma 1. *The best response threshold $t_{BR(\theta)}$ decreases when the threshold θ increases.*

Proof. We first observe that if θ increases, then the number of vaccinated population increases, which implies that the number of infected population $m_I(t)$ decreases. From the proof of Proposition 1, we know that $J_I(t) \geq J_S(t)$ for all $t \geq t_{BR(\theta)}$. Thus, in (3) $m_I(t)$ is multiplied by $J_I(t) - J_S(t)$, which is positive. Therefore, if the number of infected population $m_I(t)$ decreases then $J_S(t)$ also decreases for all $t \geq t_{BR(\theta)}$. Finally, this implies that $t_{BR(\theta)}$ decreases. □

According to Definition 1, a strategy with threshold θ is a mean-field equilibrium if $t_{BR(\theta)} = \theta$. From Lemma 1, and continuity, by letting θ increase from 0 to T , the threshold $t_{BR(\theta)}$ meets θ once, which gives the following proposition.

Proposition 2. *There exists a unique pure mean-field equilibrium that is a threshold strategy (up to a set of measure zero).*

An important consequence of this result is the simplification of the numerical computation of a mean-field equilibrium since it can be done by solving a fixed point problem over a scalar value (see 6 for actual numerical computations). In the following section, we show that the global optimum of the problem is also of threshold type.

4. Social Optimum Characterization

In this section, we consider the social optimum strategy, that minimizes the global cost of the whole population (also called the social cost).

4.1. Centralized Control Problem

We focus on a centralized control problem for this epidemic model. Here, we seek to find the vaccination strategy of the population π such that the total cost of the system is minimized. We denote by $C(\pi)$ the *social cost* incurred under the population vaccination strategy π , i.e.,

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

The optimal social cost is obtained by applying an optimal population strategy π^{opt} :

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

4.2. Social Optimum Characterization

The authors in [6] have shown that the optimal social cost is the unique viscosity solution of an Hamilton-Jacobi-Bellman equation. In particular this implies that there exist solutions that are measurable. In this section, we show that the centralized control problem is also tractable and that one can characterize its solution.

As for the case of mean-field equilibrium, a policy is essentially defined up to a set of measure 0 in the following sense: If

two policies differ on a null set (*i.e.* of measure 0), then the cost and the state remain identical.

In the following we will show that a global optimum is reached using a threshold strategy (again up to a null set).

Proposition 3. *Any strategy that minimizes the total cost is a threshold strategy.*

Proof. The proof is based on an improving policy argument. In the rest of the proof, we never refer to the fact that policies are defined up a null set.

Consider a population strategy π_0 that is not a threshold policy. This means that a small fraction of the vaccination rate can be moved to the left. More precisely, this implies that there exist $0 < t_0 < T$, such that for all $\delta > 0$, there exists a measurable function $\eta(t)$, and $\varepsilon > 0$ such that:

- $\varepsilon < \delta^3$,
- $\int_{t_0}^{t_0+\varepsilon} \eta(t)m_S(t)dt = \varepsilon$,
- $\int_{t_0+\delta}^{t_0+\delta+\varepsilon} \eta(t)m_S(t)dt = -\varepsilon$,
- $\eta(t) = 0$ outside these two intervals,
- The policy $\pi_1 := \pi_0 + \eta$ is a valid policy: $\forall 0 \leq t \leq T, \pi_0(t) + \eta(t) \in [0, \tau]$.

The policy π_1 can be seen as a small shift of vaccination rate to the left, just after time t_0 . If $m_S^1(t)$ and $m_I^1(t)$ are the proportions of susceptible and infected populations under strategy $\pi_1(t)$ and $m_S^0(t)$ and $m_I^0(t)$ be the proportions of susceptible and infected populations under strategy $\pi_0(t)$, then let us further define $u_0(t) = \pi_0(t)m_S^0(t)$ and $u_1(t) = \pi_1(t)m_S^1(t)$.

We will now show that $\pi_1(t)$ is an improving strategy compared with $\pi_0(t)$, *i.e.*

$$\int_0^T (c_I m_I^1(t) + c_V u_1(t)) dt < \int_0^T (c_I m_I^0(t) + c_V u_0(t)) dt. \quad (7)$$

We will split the integral over several intervals $[0, t_0]$, $[t_0, t_0 + \varepsilon]$, $[t_0 + \varepsilon, t_0 + \delta]$, $[t_0 + \delta, t_0 + \varepsilon + \delta]$, $[t_0 + \varepsilon + \delta, t_0 + 2\delta]$, $[t_0 + 2\delta, T]$ and compare the costs on each on them.

Let $i_0 := m_I^1(t_0) = m_I^0(t_0)$ and $s_0 := m_S^1(t_0) = m_S^0(t_0)$.

- $[0, t_0]$: Both policies are identical up to t_0 so that the costs and states coincide.
- $[t_0 + \varepsilon, t_0 + \delta]$: Using a linearization around t_0 , we obtain for any $\varepsilon < t < \delta$,

$$\begin{aligned} m_I^1(t_0 + t) &= m_I^0(t_0 + t) - t^2 \gamma \varepsilon i_0 / 2 + O(\varepsilon t^3), \\ m_S^1(t_0 + t) &= m_S^0(t_0 + t) - t \varepsilon + t^2 \gamma \varepsilon i_0 / 2 + O(\varepsilon t^3). \end{aligned}$$

If δ is small enough, then $m_I^0(t_0 + t) > m_I^1(t_0 + t)$ for all $t < \delta$.

- $[t_0 + \delta + \varepsilon, t_0 + 2\delta]$:

Using linearization again, for any $\varepsilon < t < \delta$,

$$\begin{aligned} m_S^1(t_0 + \delta + t) &= m_S^0(t_0 + \delta + t) + t^2 \varepsilon \gamma i_0 / 2 + O(\varepsilon t^3), \quad (8) \\ m_I^1(t_0 + \delta + t) &= m_I^0(t_0 + \delta + t) - t^2 \varepsilon \gamma i_0 / 2 + O(\varepsilon t^3). \quad (9) \end{aligned}$$

Again, if δ is small enough, then these equations imply dominance of π_1 over both intervals $[t_0 + \delta + \varepsilon, t_0 + 2\delta]$ and $[t_0 + \varepsilon, t_0 + \delta]$.

- $[t_0, t_0 + \varepsilon]$ and $[t_0 + \delta, t_0 + \delta + \varepsilon]$: Over both these intervals, the difference in cost between both policies is $O(\varepsilon \delta^3)$ so that it becomes negligible when δ is small enough.
- $[t_0 + 2\delta, T]$:

From time $t_0 + 2\delta$ on, both policies coincide. By monotonicity of the evolution equations, if $m_I^0(t') \geq m_I^1(t')$ for $t' > t_0 + 2\delta$, then this remains true thereafter.

Finally, we have shown that any non threshold policy can be strictly improved. Therefore, we conclude that the population strategy of minimal cost is a threshold strategy. \square

Proposition 4. *The threshold of the social optimal θ^* is always larger than the threshold of the Nash equilibrium, θ^{eq} .*

Proof. The proof is based on the Pontryagin's maximum principle applied to our optimization problem. The Pontryagin's maximum principle gives a necessary condition for optimality [12]. In our case, it translates into the following: If π^* is an optimal strategy, then there exist two Lagrangian multipliers $\lambda_S(t)$ and $\lambda_I(t)$ such that $\lambda_S(T) = 0$, $\lambda_I(T) = 0$ and for any $t < T$,

$$\begin{aligned} -\dot{\lambda}_S &= c_V \pi^*(t) + (-\gamma m_I^*(t) - \pi^*(t)) \lambda_S + \gamma m_I^*(t) \lambda_I \\ -\dot{\lambda}_I &= c_I - \gamma m_S^*(t) \lambda_S + (\gamma m_S^*(t) - \rho) \lambda_I \\ \pi^*(t) &= \arg \min [c_V \pi(t) m_S^*(t) + c_I m_I^*(t) + (\gamma m_S^*(t) m_I^*(t) \\ &\quad - \pi(t) m_S^*(t)) \lambda_S + (\gamma m_S^*(t) m_I^*(t) - \rho m_I^*(t)) \lambda_I], \end{aligned}$$

where $m_I^*(t), m_S^*(t)$ are the proportions of the population in states I and S respectively, at time t , under the optimal strategy.

By straightforward simplifications, one gets

$$-\dot{\lambda}_S = \inf_{\pi} (\pi(t)(c_V - \lambda_S) + \gamma m_I^*(t)(\lambda_I - \lambda_S)) \quad (10)$$

$$-\dot{\lambda}_I = c_I - \rho \lambda_I + \gamma m_S^*(t)(\lambda_I - \lambda_S) \quad (11)$$

$$\pi^*(t) = \arg \min (\pi(t) m_S^*(t) (c_V - \lambda_S)). \quad (12)$$

Equations (10)-(11) are similar to the equation for the costs of the best response policy (3)-(4) up to the additional term $\gamma m_S^*(t)(\lambda_I - \lambda_S)$ for λ_I . Using this, the comparison between the optimal strategy and the mean field equilibrium boils down to the comparisons of the Lagrangians λ_S, λ_I and the costs J_S, J_I .

One easy case is where c_V is larger than c_I/ρ . In this case, for all t , $J_S \leq J_I \leq c_I/\rho \leq c_V$ so that the threshold of the Nash equilibrium is $\theta^{eq} = 0$. Therefore, the social optimal threshold θ^{opt} can only be larger than θ^{eq} .

Let us now consider the case where $c_V < c_I/\rho$. In this case, θ^* is the time when λ_S gets below c_V . By examining the Lagrangian multipliers λ_S and λ_I between θ^* and T , one can show that they must satisfy the following properties:

- $\lambda_S(T) = 0, \lambda_I(T) = 0,$
- $\forall t \in [\theta^*, T], \lambda_S(t) \leq \lambda_I(t).$

Indeed, if there is a time t such that $\lambda_S(t) = \lambda_I(t)$, then their derivatives become comparable ($\dot{\lambda}_S(t) \leq \dot{\lambda}_I(t)$). Therefore, the additional term $\gamma m_S^*(t)(\lambda_I - \lambda_S)$ in (10) remains positive so that $\lambda_I(t) \leq J_I(t), \forall \theta^* \leq t \leq T$. In turn this implies that $\lambda_S(t) \leq J_S(t), \forall \theta^* \leq t \leq T$.

This implies that $J_S(\theta^*) \leq \lambda_S(\theta^*) = c_V$.

Finally, this implies that θ^{eq} (the time when J_S crosses level c_V) is smaller than θ^* . \square

5. Comparison with the N player game

If the population is made of a finite number of players (N), the existence of Nash equilibria (NE) follows from classical results on stochastic games over finite horizons. However, computing or even characterizing NE is not easy and remains open even in our simple framework.

Using Theorem 3 in [3], we can claim that the threshold MFE computed in Section 3 is an ε NE of the N player game, and ε goes to 0 when N goes to infinity.

Conversely, it might be possible that the N player game has several equilibria. In particular it might have non-symmetrical NE and/or mixed equilibria. We believe that this is not the case here. Indeed, one can exploit the fact that the game has a finite time horizon and bounded rates to construct a backward induction starting from T and moving back in time one event at a time to show that the game has a single equilibrium.

Actually, even if the game has several equilibria whose social costs are strictly better than the social cost of the MFE, these NE must have a complex information structure: They cannot be *local* in the sense of [3]. For non-local strategies, the decisions of one player will depend on the decisions of the other players. These types of NE may not be realistic in practice, because players cannot know the vaccination strategies of all the other players.

6. Numerical comparisons

Proposition 2 and Proposition 3 prove that the mean-field equilibrium and the global optimum are threshold strategies. In this section, we compute the thresholds of these strategies numerically. We consider the same system parameters² as in

²Actually, the model of [7] is very similar but their definition of the parameter c_I is different. The c_I in [7] corresponds to our value of c_I/ρ . Hence, in [7] they use $c_I = 1$ which corresponds to $c_I/\rho = 36.5$. It is easy to see that if the cost of infection is c_I times ρ both models coincide. The authors, using their approximation, obtain that the cost for the mean-field equilibrium and in the global optimum are, respectively, 0.55 and 0.53. Using our approach, the genuine costs are 0.542 for the mean-field equilibrium and 0.524 for the global optimum.

[7], which is based on the epidemiological study of the H1N1 epidemic of 2009-2010 in France : $\rho = 36.5, \gamma = 73, \tau = 10, T = 0.3, c_I = 36.5$ and $c_V = 0.5$.

We computed the optimal strategy and the mean-field equilibrium. The results are reported in Figure 2a where the population state space is divided in three regions that represent the decisions taken by both policies at time 0, as a function of the initial state. In the white region, both strategies vaccinate at maximum rate. In the dark gray region, the strategy of the global optimum is to vaccinate at maximum rate and the strategy of the equilibrium is to not vaccinate. In the light gray region, both strategies are to not vaccinate.

We also plot the trajectories corresponding to both strategies when the initial proportion of infected population and of susceptible population are both equal to 0.4 at time 0. In Figure 2a (see Figure 2b for a zoomed figure), we plot with a solid line the behavior of the equilibrium vaccination strategy, and with a dashed line, the behavior of the global optimum.

For any vaccination cost c_V , while the other parameters remain fixed, we denote by $t^{opt}(c_V)$ (resp. $t^{eq}(c_V)$) the threshold of the global optimum strategy (resp. equilibrium strategy). It was shown that in both cases, the thresholds are decreasing in c_V : the more costly is the vaccination, the less people vaccinate (for the globally optimal situation or for the mean-field equilibrium).

Figure 3 confirms that the thresholds decrease with c_V and also shows that the thresholds are never equal for this range of parameters. This suggests that, if the vaccination decisions are let to individuals, then vaccination should be subsidized, by removing a grant g to the vaccination cost in the equilibrium so that both thresholds coincide, *i.e.*,

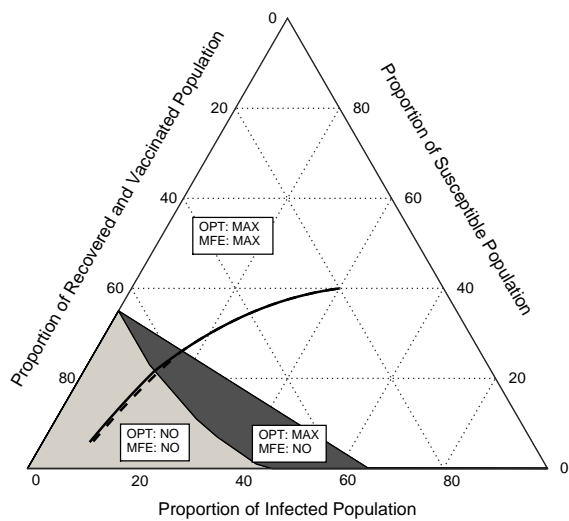
$$t^{eq}(c_V - g) = t^{opt}(c_V).$$

For example, with the same parameters as in the simulation of Figure 2, we observe that, for $c_V = 0.8$, the threshold of the global optimum is 0.034, while the threshold of the equilibrium is 0. As it can be seen, the threshold of the equilibrium is 0.034 when $c_V = 0.45$. This simulation shows that the subsidy required to encourage selfish individuals to vaccinate optimally consists of a reduction of the vaccination cost of $g = 0.35$.

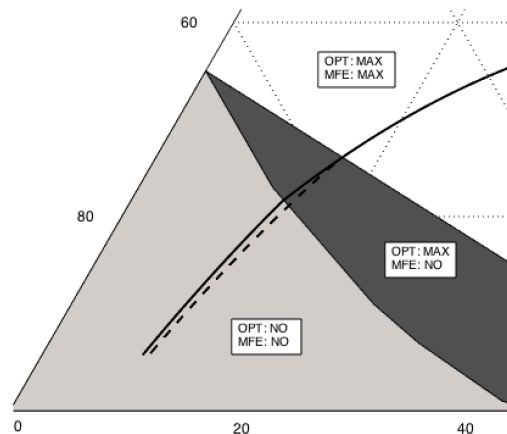
7. Conclusions

This study has an interest per se : exhibiting the gap between social optimum and mean-field equilibrium in a specific case. Moreover, it can be seen as a proof of concept showing that mean-field equilibrium can be computed under almost closed form when the state space of each player is finite while the N -player Nash equilibrium is impossible to compute.

Also, this is an example where the mean-field equilibrium is unique and pure. Actually, the condition for uniqueness and pureness of the equilibrium boils down to the uniqueness of the solution of the Bellman optimality equations (here Equations (3) - (4)). This brings to the natural question of finding general conditions insuring the existence of pure equilibria in mean-field games.



(a) Population dynamics.



(b) Population dynamics (zoomed).

Figure 2: Population dynamics under the equilibrium strategy (dashed line) and the global optimum strategy (solid line). Three zones are displayed: (i) in the white region, the global optimum and the equilibrium vaccinate with maximum rate; (ii) in the dark gray region, the global optimum vaccinates with maximum rate, while the equilibrium does not vaccinate; and (iii) in the light gray region, neither the global optimum nor the equilibrium vaccinates. $m_I(0) = m_S(0) = 0.4$.




Figure 3: Comparison of the threshold of the equilibrium and of the global optimum when c_V varies from 0.01 to 1.21.

References

- [1] A. Bensoussan, J. Frehse, and P. Yam. *Mean field games and mean field type control theory*. Springer, 2013.
- [2] R. Carmona and F. Delarue. Probabilistic analysis of mean-field games. *SIAM Journal on Control and Optimization*, 51(4):2705–2734, 2013.
- [3] J. Doncel, N. Gast, and B. Gaujal. Mean-Field Games with Explicit Interactions. Technical report, 2016. <https://hal.inria.fr/hal-01277098>.
- [4] O. Guéant, J.-M. Lasry, and P.-L. Lions. Mean field games and applications. In *Paris-Princeton Lectures on Mathematical Finance 2010*, volume 2003 of *Lecture Notes in Mathematics*, pages 205–266. Springer Berlin Heidelberg, 2011.
- [5] T. G. Kurtz. *Approximation of population processes*, volume 36. SIAM, 1981.
- [6] L. Laguzet and G. Turinici. Global optimal vaccination in the sir model: properties of the value function and application to cost-effectiveness analysis. *Mathematical biosciences*, 263:180–197, 2015.
- [7] L. Laguzet and G. Turinici. Individual vaccination as nash equilibrium in a sir model with application to the 2009–2010 influenza a (h1n1) epidemic in france. *Bulletin of Mathematical Biology*, 77(10):1955–1984, 2015.
- [8] J.-M. Lasry and P.-L. Lions. Jeux à champ moyen. i–le cas stationnaire. *Comptes Rendus Mathématique*, 343(9):619–625, 2006.
- [9] J.-M. Lasry and P.-L. Lions. Jeux à champ moyen. ii–horizon fini et

contrôle optimal. *Comptes Rendus Mathématique*, 343(10):679–684, 2006.

- [10] J.-M. Lasry and P.-L. Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.
- [11] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [12] E. Todorov. Optimal control theory. *Bayesian brain: probabilistic approaches to neural coding*, pages 269–298, 2006.