

A Mean Field Game Analysis of SIR Dynamics with Vaccination

Josu Doncel*, Nicolas Gast and Bruno Gaujal †

October 6, 2020

Abstract

We analyze a mean field game model of SIR dynamics (Susceptible, Infected, Recovered) where players choose when to vaccinate. We show that this game admits a unique mean field equilibrium that consists in vaccinating at maximal rate until a given time and then not vaccinating. The vaccination strategy that minimizes the total cost has the same structure as the mean field equilibrium. We prove that the vaccination period of the mean field equilibrium is always smaller than the one minimizing the total cost. This implies that, to encourage optimal vaccination behavior, vaccination should always be subsidized. Finally, we provide numerical experiments to study the convergence of the equilibrium when the system is composed by a finite number of agents (N) to the mean field equilibrium. These experiments show that the convergence rate of the cost is $1/N$ and the convergence of the switching curve is monotone.

1 Introduction

We consider a system in which the dynamics of the population is given by the classical SIR (Susceptible-Infected-Recovered) model [19] to which we add a control action for each player: each individual player can choose how and when she gets vaccinated. Each individual player tries to minimize a cost composed of a vaccination cost plus a cost function of the duration of infection. Potential applications include controlling disease spreading in a large population, anti-virus usage in computer systems or limiting the propagation of advertisements.

In this paper, we model the problem of selfish vaccination in the SIR model as a *finite state mean field game* in continuous time. Because all players have a finite state space, the state evolution will become a Kolmogorov equation and the best response strategy will be given by a Bellman optimality equation. Here, the state space only has three states, and the dynamics is simple enough so that

*J. Doncel is with the University of the Basque Country, Leioa 48940, Spain (email: josu.doncel@ehu.eus)

†N. Gast and B. Gaujal are with INRIA and Univ. Grenoble Alpes, CNRS, LIG, F-38000 Grenoble, France (email: nicolas.gast@inria.fr, bruno.gaujal@inria.fr)

one can show that this Bellman optimality equation has a unique solution. We show that the game admits a unique and pure equilibrium. Furthermore, we show that this unique equilibrium consists in a bang-bang strategy : each player chooses to vaccinate at maximal rate until some critical time τ^{eq} and then stops vaccinating thereafter.

We also compare the global cost of the mean field equilibrium to the cost of a centralized optimal control. It is known that the solution of this problem is also a bang-bang strategy [4]: it consists in vaccinating at maximum rate until a given jump time τ^* is reached and to stop vaccinating after this. We show that the time at which the global optimum stops vaccinating is always larger than that of the mean field equilibrium : $\tau^{eq} \leq \tau^*$, with equality only in degenerate cases. This shows that if the decision of whether to vaccinate or not is left to individual, then vaccination should be subsidized in order to reduce the total cost of the mean field equilibrium down to the one of the socially optimal strategy. Similar results have been obtained on SIR models with vaccination (see the related work section).

One of the main contribution of our paper is to recast the problem of selfish vaccination as a finite state mean field game (FMFG). This allows to

- Use the general theory of FMFG to assert the existence of an equilibrium;
- Use a generic approach to solve the problem, *i.e.* instantiate the coupled equations of a FMFG (Kolmogorov + Bellman) in the current case (this is done in §4), and solve them symbolically or numerically §4 and §5);

We also assess the relations between the mean field case and the case with a finite number of players. In Section 6, we present how the equilibrium strategy of the N -player game, as well as the social optimum, can be computed using the Bellman equation. Moreover, in Section 7, we study numerically the convergence of the case with N players to the mean field case. Our numerical experiments show that, for the equilibrium and for the social optimum, the convergence rate of the cost is $1/N$ and the convergence of the switching curve is monotone. Finally, we study the efficiency of the Nash equilibrium with a small number of players and of the mean field equilibrium.

The rest of the paper is organized as follows. In Section 2, we introduce the FMFG model of the problem. In Section 3, we describe the model and we define the notion of mean field equilibrium. We characterize the equilibrium of the mean field game in Section 4. We analyze the centralized control problem in Section 5. We comment on the difference between the mean field equilibrium and the Nash equilibrium of a game with a finite number of players in Section 6. We compare numerically the mean field equilibrium and the global optimum in Section 7. Finally, we conclude in Section 8.

2 Related Work

The SIR dynamics is one of the simplest and one of the most studied model of virus propagation in a moving population or information diffusion in social net-

works, since the work of Kermack and McKendrick [19]. The problem of optimal vaccination that allows the population to be immune to the infection has been analyzed in several papers, see for example [1, 9] for monographs in this topic. The authors in [21] show that the optimal social cost is the unique viscosity solution of an Hamilton-Jacobi-Bellman equation. Using the Pontryagin's Maximum Principle, the authors in [4] characterize the vaccination strategy that minimizes the total cost of a model that covers most of the literature regarding deterministic control of SIR epidemics, see [23] for a particular case of the model of [4] where $f(x, y) = xy$, g, c constant and $T \rightarrow \infty$. As we will see below, the centralized control problem we consider is also a particular case of the model of [4] and, therefore, we can use their result to study the compare with the efficiency of individual vaccinations. The authors in [14] and in [2] also perform this comparison considering, respectively, the disease eradication and the SIR dynamics in steady state.

In this paper, we consider a game scenario in which each player chooses when and how to vaccinate. On this topic, the paper closest to ours is [11] that studies a model aiming to minimize the cost of infection and of vaccination in a SIR model within a given population. Similarly to our case, the authors of [11] show that an equilibrium can be characterized by a switching curve, which is the set of points such that individuals see no difference between vaccinating and not doing when no one else vaccinates. They prove the uniqueness of such a curve and conclude thus that individual vaccinations consists of vaccinating with maximum rate until the switching curve is reached and then not to vaccinate again. One of the main contribution of our work is to recast the model of [11] in the more precise context of mean field game and show that an equilibrium under their definition corresponds to a fixed point of the best response correspondence for each player (MFG were not yet defined when [11] was published).

Our model is very close to the one of [22] in which the authors consider a mean field game model of vaccination, but with a few differences. First, the authors of [22] model the problem following closely the notion of timing games whereas we formulate our model directly as a discrete-space mean field games. The authors in [18] extend the results of [22] considering additional birth and death rates on the population. We believe that our formulation is simple and serves as a good illustration of the expressive power and simplicity of discrete-space mean field games, as our formulation is not specific to the vaccination problem and could be easily generalized. Besides, as in [11], the vaccination policies in [22] do not depend explicitly on the susceptible population whereas our rate does: Essentially, the dynamics of our Equation (1) includes a term $\pi(t)m_s(t)$ whereas the one of [22] has only the term $\pi(t)$. Using a term $\pi(t)m_s(t)$ makes our model fall in the category of discrete-space mean field game of [10]. In terms of results, we see two main contributions of our paper with respect to [22]. First, we provide alternative proofs to the results of [22] that are based on structural properties of the discrete-space mean field game and that rely less on the explicit form of the dynamics. Second, we provide a detailed comparison with the finite-player model.

Mean field games with discrete and finite state space have been explored in

discrete time ([10, 15]) as well as in continuous time ([16, 17]). Our model is a particular case of a mean field game with finite state space evolving in continuous time (see [10]). The existence of an equilibrium in our mean field game follows directly from their result. Here, however, we go beyond this existence result and we show that there is a unique mean field equilibrium that happens to be a deterministic strategy and we characterize it.

3 Model Description

3.1 System Dynamics

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) or 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, in state S , in state I or in state R 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 at rate γ . If a player is susceptible and encounters an infected player, then she becomes infected. An infected player recovers at rate ρ . We also consider that a susceptible player can choose a vaccination strategy π . A strategy π is a measurable function from $[0, T]$ to $[0, \theta]$, where $\theta < \infty$ is the maximal vaccination rate and $\pi(t)$ is the rate at which the player becomes vaccinated. Once a player is vaccinated or has recovered, her state does not change. The Markovian behavior of a player is represented in Figure 1.




Figure 1: The dynamics of a player in the epidemic model. A 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 sum of two costs :

- 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 with an infinite number of players. In this case, the dynamics of the population is given by the Kolmogorov Equation that takes the form of a system of ordinary differential

equations, the state space being finite.

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

If the strategy $\pi(t)$ is not continuous in time, the differential equation (1) may not be well-defined at time-points where π is not continuous. The existence of a continuous solution for (1) is guaranteed by the Carathéodory's Existence Theorem. This solution is essentially unique (*i.e.* unique up to a set of measure zero). It is the essentially unique fixed point of the following integral equation:

$$\begin{cases} m_S(t) = m_S(0) + \int_0^t (-\gamma m_S(u)m_I(u) - \pi(u)m_S(u))du \\ m_I(t) = m_I(0) + \int_0^t (\gamma m_S(u)m_I(u) - \rho m_I(u))du \\ m_R(t) = m_R(0) + \int_0^t (\rho m_I(u) + \pi(u)m_S(u))du. \end{cases}$$

The same is true for all differential equations used in the paper, without further mention to this point: all solutions can be considered continuous everywhere and differentiable at all points except over a set of measure 0.

From this equation, we note that the vaccination strategy we consider can be interpreted as a time varying parameter, whereas the authors in [22] consider a general function that do not depend explicitly on the susceptible population. Indeed the dynamics of the system of the mean field game analyzed in [22] coincide with (1) if we replace $\pi(t)m_S(t)$ by $dU(t)$.

3.2 Best Response

We focus on a particular player, that we call Player 0. As we consider a mean field game model, the dynamics of the global population is not affected by Player 0 alone and is driven by Equation (1). Player 0 chooses her vaccination strategy π^0 , where $\pi^0(t) \in [0, \theta]$ for all t . The probability that Player 0 is in a given state depends not only on π^0 , but also on the population distribution, which is the vector $m(t) = (m_S(t), m_I(t), m_R(t))$.

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:

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

Note that the above equation is similar to Equation (1) except that it is linear in x whereas Equation (1) is not linear in m . We will exploit this linearity in Section 4 to show that the optimal strategy can be computed easily by solving the Bellman's equation.

The expected individual cost of Player 0 over interval $[t, T]$ is:

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

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,

$$BR(\pi) \in \arg \min_{\pi^0} V^{\pi^0, \pi}, \quad (2)$$

which is a non-empty set by compactness of the strategy space¹.

As in [10, 15] we define a mean field equilibrium as a fixed point of the best-response function:

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

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

This is the classical definition of an equilibrium in a mean field game. The rationale behind this definition is that in a homogeneous population, each player's best-response is the same as for Player 0. For a given vaccination strategy of the population π , any player of the population choose the vaccination strategy $BR(\pi)$. As in classical games, a mean field equilibrium is a situation where no player has incentive to deviate unilaterally from the selected vaccination strategy.

4 Mean Field Equilibrium Characterization

In this section, we analyze the mean field equilibrium described in Section 3.2. 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, our model is a particular case of the finite mean field games whose equilibrium existence is proven in [10]. In this section, we use a direct method to further characterize an equilibrium.

We define a bang-bang policy with at most one jump the policy that vaccinates at maximum rate until a given jump time τ and does not vaccinate after time τ . We denote this policy by π^τ . More precisely,

$$\pi^\tau(t) = \begin{cases} \theta & \text{if } t < \tau, \\ 0 & \text{if } t \geq \tau. \end{cases}$$

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

In this section we show that the mean field equilibrium is unique, deterministic (or pure), in the sense that players do not randomize when the equilibrium is reached, and is a bang-bang strategy that jumps at a time that we denote by τ^{eq} .

The first step in our proof is to show that, for any strategy π (not necessarily bang-bang) there exists a unique best response to π that is bang-bang strategy with at most one jump.

Proposition 1. *For any population strategy π , the best-response to π is unique, up to a set of measure 0. It is a bang-bang strategy with at most one jump. The jump time is denoted by $\tau_{BR}(\pi)$.*

Proof. From the discussion in Section 3.2, 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 best response to π , satisfy the Hamilton Jacobi Bellman equations [5]:

$$J_S(T) = J_I(T) = 0 \quad (3)$$

$$-\dot{J}_S(t) = \inf_{\pi^0} [\pi^0(t)(c_V - J_S(t)) + \gamma m_I(t)(J_I(t) - J_S(t))] \quad (4)$$

$$-\dot{J}_I(t) = c_I - \rho J_I(t). \quad (5)$$

$$\begin{aligned} BR(\pi) &= \arg \min_{\pi^0} [\pi^0(t)(c_V - J_S(t)) + \gamma m_I(t)(J_I(t) - J_S(t))] \\ &= \arg \min_{\pi^0} [\pi^0(t)(c_V - J_S(t))] \end{aligned} \quad (6)$$

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.

The solution of Equation (3)-(5) is $J_I(t) = \frac{c_I}{\rho}(1 - e^{\rho(t-T)})$. Therefore, the cost $J_I(t)$ is decreasing from $\frac{c_I}{\rho}(1 - e^{-\rho T})$ to 0. At time T , $J_S(T) = 0 = J_I(T)$ so that $\dot{J}_S(T) = 0$. Also, by (4), $-\dot{J}_S(t) \leq \gamma m_I(t)(J_I(t) - J_S(t))$. By continuity of the functions J_S and J_I , this implies that $J_S(t) < J_I(t)$ for all $t \in [0, T)$.

By (6), the best response strategy to π is 0 if $J_S(t) < c_V$ and θ 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 $J_S(t) = c_V$ (otherwise, we set $t_1 = 0$). Equation (4), together with the fact that $J_S(t) < J_I(t)$ implies that $-\dot{J}_S(t_1) > 0$. By continuity, for all $t > t_1$, one has $-\dot{J}_S(t) > 0$. This implies that J_S is decreasing from t_1 to T . Hence, $J_S(t)$ crosses level c_V at most once (at a time t_1). This implies that the best response to π is a bang-bang strategy with at most one jump:

$$BR(\pi) = \begin{cases} \theta & \text{if } t < t_1 \\ 0 & \text{if } t \geq t_1. \end{cases} \quad (7)$$

□

Notice that the best response is not necessarily well-defined at the jump time. This is in accordance with the fact that two strategies that differ up to a set of measure 0 will lead to the same costs. Another consequence of the previous proof is that the jump time $BR(\pi)$ is “monotonous” in π in a sense precised in the following result.

Lemma 1. *Let π^τ be a bang-bang strategy that jumps at τ . Then, the jump time $\tau_{BR(\pi^\tau)}$ of the response to π^τ decreases when the original jump time τ increases.*

Proof. We first observe that if τ 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 $t \geq \tau_{BR(\pi^\tau)}$. Thus, in (4) $m_I(t)$ 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 $t \geq \tau_{BR(\pi^\tau)}$. This implies that $\tau_{BR(\pi^\tau)}$ decreases. \square

The population distribution is a continuous function of the threshold τ and, as a consequence, so are $J_S(t)$ and $J_I(t)$. This implies that the mapping $(\tau^0, \tau) \mapsto V^{\pi^{\tau^0}, \pi^\tau}$ is also continuous. By Proposition 1, the best response $\tau_{BR(\pi^\tau)} = \arg \min_{\tau^0} V^{\pi^{\tau^0}, \pi^\tau}$ is unique. This implies that the mapping $\tau \rightarrow \tau_{BR(\pi^\tau)}$ is continuous. From Lemma 1, and by letting τ increase from 0 to T , the jump time $\tau_{BR(\pi^\tau)}$ is equal to τ exactly once. This proves the following proposition.

Proposition 2. *There exists a unique pure mean field equilibrium and it is a bang-bang strategy with at most one jump.*

An important consequence of this result is that it provides a simple way of computing numerically the mean field equilibrium as it can be done by solving a fixed point problem over a scalar value. Notice also that it is in accordance with the results obtained in [22] and [11]: the vaccination of individuals have the same structure as the mean field equilibrium we have presented.

In the following section, we present that the global optimum of the problem is also a bang-bang strategy with at most one jump and we compare the jump times of both strategies.

5 Social Optimum Characterization

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

5.1 Centralized Control Problem

Let us 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 control strategy that minimizes $C(\pi)$ among all the measurable functions from $[0, T]$ to $[0, \theta]$ is denoted by π^{opt} and satisfies

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

As for the case of the mean field equilibrium, a strategy is essentially defined up to a set of measure 0 in the following sense: If two strategies differ on a null set (*i.e.* of measure 0), then the cost and the state remain identical.

5.2 Social Optimum Characterization

In this section, we study the vaccination strategy that minimizes the total cost of the SIR model. The authors in [4] consider a model that covers most of the literature regarding deterministic vaccination strategies of SIR dynamics and show the vaccination strategy that minimizes the cost of the system a bang-bang strategy with at most one jump. We note that the centralized control problem we present in Section 5.1 is a particular case of their result, where $f(S(t), I(t)) = \beta S(t)I(t)$, $a = d = 0$ and $c(S(t)) = S(t)c_V/c_I$. Therefore, their result can be used to conclude that the social optimum of the SIR model is reached by a bang-bang strategy with at most one jump. In the rest of the section, the jump time of the social optimum will be denoted by τ^* .

Proposition 3. [4, Theorem 2.1.] *There is a unique strategy that minimizes the total cost of the population and it is a bang-bang strategy, i.e., $\pi_{opt} = \pi^{\tau^*}$.*

In the following result, we show that, unless for trivial cases, the mean field equilibrium and the social optimum never coincide. More interestingly, we also prove that the jump time of the mean field equilibrium τ^{eq} is always smaller than the optimal jump time τ^* . This implies that rational individuals who try to minimize their own cost would tend to be less vaccinated than what a centralized controller would do, essentially because they benefit from the vaccination of the others. Hence, in order to encourage individuals to vaccinate optimally, a central planner should subsidize vaccination.

Proposition 4. *The jump time τ^* of the socially optimal strategy is always equal or larger than the jump time of the mean field equilibrium τ^{eq} , with equality only if $\tau^* = 0$.*

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 [25]. In our case, it translates into the following: If

π^* is an optimal strategy, then there exist two Lagrange multipliers $\lambda_S(t)$ and $\lambda_I(t)$ such that $\lambda_S(T) = 0$, $\lambda_I(T) = 0$ and for any $t < T$,

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

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

$$-\dot{\lambda}_S = \inf_{\pi} (\pi(t)(c_V - \lambda_S) + \gamma m_I^*(t)(\lambda_I - \lambda_S)) \quad (8)$$

$$-\dot{\lambda}_I = c_I - \rho\lambda_I + \gamma m_S^*(t)(\lambda_I - \lambda_S) \quad (9)$$

$$\pi^*(t) = \arg \min (\pi(t)m_S^*(t)(c_V - \lambda_S)). \quad (10)$$

Equations (8)-(9) are similar to the equation for the costs of the best response strategy (4)-(5) 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 Lagrange Multipliers 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 , $\lambda_S \leq \lambda_I \leq c_I/\rho \leq c_V$ so that the jump time of the mean field equilibrium is $\tau^{eq} = 0$. Therefore, the socially optimal jump time τ^* can only be larger than τ^{eq} .

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

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

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 (8) remains positive so that $\lambda_I(t) \leq J_I(t), \forall \tau^* \leq t \leq T$. In turn this implies that $\lambda_S(t) \geq J_S(t), \forall \tau^* \leq t \leq T$.

This implies that $J_S(\tau^*) \leq \lambda_S(\tau^*) = c_V$. Finally, this implies that τ^{eq} (the time when J_S crosses level c_V) is smaller than τ^* (with equality only when $\tau^* = 0$). \square

This result shows the difference between the vaccination of the individuals when they behave in a selfish manner and the vaccination that minimizes the cost of the SIR model. Indeed, we have proven that the jump time of the social optimum is always larger than the jump time of the mean field equilibrium. In Section 7, we compare both vaccination policies numerically (see [6] for a related

comparison between the forward-backward stochastic differential equations of a mean field game and the centralized control problem over controlled dynamics of McKean-Vlasov type). Prior to that, we study this SIR model with finite number of players in the following section.

6 Analysis with $N < \infty$ players

An alternative model would have been to consider a population made of a finite number of players (N). In this section, we present vaccination models with $N < \infty$ players. We give the Bellman equation for this model and we discuss how it can be used to compute the Nash equilibrium and the social optimum. We first consider the Nash equilibrium and then we focus on the social optimum.

6.1 Nash Equilibrium

We consider a population formed by a finite number of players, that interact through a *stochastic game*. The existence of Nash equilibria for such a model follows from classical results on stochastic games [24].

The state of the $N + 1$ players (N players plus Player 0) can be seen as a triplet (X, M_S, M_I) where X is the state of Player 0: $X \in \{S, I, R\}$, M_S (resp. M_I) is the number of the other players in state S (resp. I).

When all the other players use strategy π , the best response of Player 0 is given by the following Bellman equation (after uniformization of the rates) over $T(N + 1)$ time steps:

$$\begin{aligned} C_t(S, M_S, M_I) &= \min_{\pi_t^0(M_S, M_I) \in [0, \theta]} p_V c_V + p_I C_{t+1}(I, M_S, M_I) + q_V C_{t+1}(S, M_S - 1, M_I) \\ &\quad + q_R C_{t+1}(S, M_S, M_I - 1) + q_I C_{t+1}(S, M_S - 1, M_I + 1) \\ &\quad + \bar{p}_S C_{t+1}(S, M_S, M_I), \\ C_t(I, M_S, M_I) &= c_I + q_V C_{t+1}(I, M_S - 1, M_I) + q_R C_{t+1}(I, M_S, M_I - 1) \\ &\quad + q'_I C_{t+1}(I, M_S - 1, M_I + 1) + \bar{p}_I C_{t+1}(I, M_S, M_I), \end{aligned}$$

where the probabilities are

$$\begin{aligned} p_V &= \Omega \pi_t^0(M_S, M_I), \\ p_I &= \Omega \gamma M_I / N, \\ p_R &= \Omega \rho, \\ q_V &= \Omega \pi_t(M_S, M_I) M_S, \\ q_R &= \Omega \rho M_I, \\ q_I &= \Omega \gamma M_S M_I / N, \\ q'_I &= \Omega \gamma M_S (M_I + 1) / N, \\ \bar{p}_S &= 1 - p_V - p_I - q_V - q_R - q_I, \\ \bar{p}_I &= 1 - p_R - q_V - q_R - q'_I, \end{aligned}$$

with a normalization $\Omega = ((N + 1)(\theta + \rho + \gamma))^{-1}$.

Using [10, Theorem 3], we can claim that the MFE computed in Section 4 is an ε -Nash Equilibrium of the N player game for any N , and ε goes to 0 when N goes to infinity. We observe that the Bellman equations above can be used to compute the best response of Player 0 to the population strategy if we replace min by argmin in the first equation. This process can be carried out iteratively to compute its fixed point (*i.e.* π^0 is the best response to itself), *i.e.*, the Nash equilibrium policy. We have not been able to prove convergence of this best response dynamics nor to prove the uniqueness of its fixed point. However, in Section 7.2, we study the Nash equilibrium for a finite number of players with the parameters of [22] and, in all the experiments we have performed, the iteration always converged to a unique point. Besides, when $N = 5$, the number of iterations required to converge is two, but for larger values of N the number of iterations increases substantially. For instance, when $N = 20$, the convergence is given after more than 800 iterations and when $N = 30$ after more than 2200 iterations.

6.2 Social optimum

We now analyze the social optimum with N players. The state of the system can be seen as a couple (M_S, M_I) where M_S (resp M_I) is the number of the other players in state S (resp. I). We aim to find the strategy that minimizes the expected cost of the system. This can be done by solving the following Bellman equation (after uniformization of the rates) over $T(N)$ time step:

$$\begin{aligned} J_t(M_S, M_I) = & c_I M_I + \min_{\pi_t \in [0, \theta]} (c_V q_V M_S + q_V J_{t+1}(M_S - 1, M_I) \\ & + q_R J_{t+1}(M_S, M_I - 1) + q_I J_{t+1}(M_S - 1, M_I + 1)) \\ & + \bar{q} J_{t+1}(M_S, M_I), \end{aligned}$$

where the probabilities are

$$\begin{aligned} q_V &= \Omega \pi_t M_S, \\ q_R &= \Omega \rho M_I, \\ q_I &= \Omega \gamma M_S M_I / N, \\ \bar{q} &= 1 - q_V - q_R - q_I, \end{aligned}$$

with a normalization $\Omega = (N(\theta + \rho + \gamma))^{-1}$.

We observe that the computation of the strategy requires a single backward computation of the Bellman equation. Note that while the socially optimal strategy is not necessarily unique, its cost $J_t(M_S, M_I)$ is uniquely defined. Moreover, its computation by the above backward recursion is much faster than the computation of the Nash equilibrium we presented in Section 6.1 because one needs only one pass of the backward recursion to obtain $J_0(M_S, M_I)$. In Section 7.2, we present numerical experiments that study the speed of convergence of the global optimum strategy with N players to the mean field optimum strategy for the parameters of [22].

7 Numerical Comparisons

7.1 Comparison of MFE and Social Optimum

Proposition 2 and Proposition 3 show that the mean field equilibrium and the social optimum are characterized by their jump time τ . In this section, we report a numerical evaluation of the jump time of these strategies². We consider the same system parameters as in [22], which is based on the epidemiological study of the H1N1 epidemic of 2009-2010 in France : $\rho = 36.5$, $\gamma = 73$, $\theta = 10$, $c_I = 36.5$ and $c_V = 0.5$. Besides, we consider that the proportion of susceptible and infected population at time 0 are both equal to 0.4.

Remark 1. *The value of c_I in [22] corresponds to our value of c_I/ρ . Hence, in [22] 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 the Bellman equations of both models coincide. The authors provide simulations where the proportion of susceptible and infected population at time 0 is, respectively, 0.75 and 0.1 and, using their approach, they obtain that the cost for the mean field equilibrium and in the social optimum are, respectively, 0.55 and 0.53. Using our approach, the costs are 0.542 for the mean field equilibrium and 0.524 for the social optimum.*




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

For these parameters, we compute the optimal strategy and the mean field equilibrium over a time horizon of a year, that is, $T = 1$. The results are reported

²The codes to reproduce these experiments are available at <https://github.com/josudoncel/MeanFieldGameAnalysisSIRModelVaccinations>

in Figure 2 where the population state space is divided into three regions that represent the decisions taken by both strategies 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 social 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 proportion of infected population and of susceptible population at time 0 are both equal to 0.4. In the left plot of Figure 2 (see the right plot 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 social optimum. The obtained cost for the equilibrium vaccination strategy for the parameters under consideration is 0.6824, whereas for the social optimum vaccination strategy is 0.6818.

For any vaccination cost c_V , while the other parameters remain fixed, we denote by $\tau^*(c_V)$ (resp. $\tau^{eq}(c_V)$) the jump time of the socially optimal strategy (resp. equilibrium strategy). It can be shown that in both cases, the jump times are decreasing in c_V : the more costly is the vaccination, the less people vaccinate (for the socially optimal situation as well as for the mean field equilibrium). Figure 3 confirms that the jump times decrease with c_V and also shows that the jump times 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 offering a grant g off the vaccination cost so that both jump times coincide, *i.e.*,

$$\tau^{eq}(c_V - g) = \tau^*(c_V).$$




Figure 3: Jump times comparison when c_V varies from 0.01 to 1. The jump time of the mean field equilibrium (MFE) is represented with a dotted line and the jump time of the social optimum (OPT) with a solid line. The horizontal distance is the subsidy to be granted to incentive players to use the optimal vaccination strategy.

For example, with the same parameters as in the simulation of Figure 2, and for $c_V = 0.8$, the jump time of the social optimum is 0.0106, while the jump

time of the equilibrium is 0. The jump time of the equilibrium is 0.034 when $c_V = 0.65$. This simulation shows that the subsidy required to encourage selfish individuals to vaccinate optimally consists of a reduction of the vaccination cost of $g = 0.15$.

7.2 Convergence of the N -player Nash Equilibrium to MFE

Using the same parameters as in the previous subsection, we compute the cost of the Nash equilibrium with N players and compare with the cost of the MFE, and let N grow to assess the convergence of the cost as N goes to infinity.

We first focus on the cost of the equilibria. As it can be observed in Figure 4, the cost of the equilibrium of the N -player game converges to the cost of the mean field equilibrium when N is large. This illustration suggests that convergence rate of the cost is in $1/N$. This rate of convergence is in accordance with the recent results of [3, 8] that show that for some specific games, the sequence of equilibria of the N -player games converges at rate $1/N$ to the mean field equilibrium³. This convergence rate is faster than the classical convergence rate in $1/\sqrt{N}$ that holds for general mean field games (see for example [7], that show that a mean field equilibrium is an $1/\sqrt{N}$ -Nash equilibrium). We believe that there are two reasons for that. First, the dynamics at equilibrium is piecewise smooth, which means that techniques similar to the ones of [12, 13, 20] could most certainly be used to prove the $1/N$ rate of convergence. Second, we focus here on the convergence of the cost of the equilibrium which tends to converge faster than the policy.




Figure 4: Cost of the Nash equilibrium with N players as N grows, and best fit of the form $a + b/N$

Figure 5 shows how the Nash policy (a switch curve for any N) converges

³Note that their results do not apply as is to our model because our model does not satisfy all the assumptions of either [3] or [8]: in those papers the player controls directly the jump rates. Moreover, those papers assume a strict convexity of costs whereas our cost is linear.

to the MFE policy. Notice that the convergence is monotone (switching curve decreases to the limit curve). Note that the computation of the mean-field equilibrium is much faster than the N -player equilibrium. In fact, we have performed these experiments in a machine equipped with a 2.7 GHz CPU core and 8 GB RAM and the computation time of the N -player equilibrium was more than 2 hours for $N = 20$ and more than 100 hours for $N = 50$, whereas the mean-field equilibrium has been computed in less than three minutes.




Figure 5: Convergence of the Nash switching curve to to the MFE switching curve

In the next set of experiments, we consider the social optimum policy. We first show in Figure 6 that the cost of the social optimum policy with a finite number of agents converges to the cost of the mean field optimum. We also remark that the convergence rate of the cost is again in $1/N$.

We compare the switching curve of the social optimum policy with N players with that of the mean field optimum policy in Figure 7. We observe from this plot that the convergence is also monotone in this case.

Note that here numerical experiments can be carried for larger values of N because the computation of the social optimal is much faster (a single backward computation of the Bellman equation) while computing the Nash equilibria requires many backward computations to reach a fixed point.

7.3 Efficiency of the N -player equilibrium

For a given value of the parameters of the model, the efficiency ratio of an equilibrium is defined as the ratio of its cost and the cost of the global optimum strategy. When the efficiency ratio is close to one, we say that the N -players equilibrium is efficient. Using this definition, the Price of Anarchy is defined as the maximum over all the parameters of the system of the efficiency ratio. From the arguments of the previous section, we know that, when the number of




Figure 6: Cost of the social optimum with N players as N grows, and best fit of the form $a + b/N$




Figure 7: Convergence of the social optimum switching curve for N players to the mean field social optimum switching curve

players is large, the efficiency ratio of an Nash equilibrium is well approximated by the mean-field efficiency ratio, i.e., the ratio of the cost of the mean-field equilibrium over the cost of the mean-field optimum.

According to the results of Figure 4 and Figure 6, one can conclude that for the parameters of [22] when the initial proportion of susceptible and infected is $m_S(0) = m_I(0) = 0.4$, the equilibrium is efficient for all N . In fact, when $N = 10$ and $N = 20$, the ratio of the cost at the equilibrium over the cost of the global optimum strategy is 1.0188 and 1.0105, respectively. Besides, for this instance, the mean-field efficiency ratio, whose value is 1.0015, approximates not that well the efficiency ratio of the N -players equilibria. In this section, we aim to study the efficiency of the equilibrium with a small number of players.




Figure 8: Ratio between the cost of the N -player equilibrium and the optimum social cost for the models with N players. We plot the ratio as a function of the initial state $(m_S(0), m_I(0))$ and compare $N = 10$ (left figure) and $N = 20$ (right figure).

We first study the efficiency of the equilibrium with a finite number of players $N = 10$ or $N = 20$. In Figure 8, we represent the efficiency ratio for all the possible initial conditions of the proportion of susceptible and infected players. As it can be observed, the efficiency ratio is close to one in most cases. In fact, for $N = 10$, the maximum of the efficiency ratio is 1.072 and it is achieved when $m_S(0) = 0.6$ and $m_I(0) = 0.1$. For $N = 20$, the cost of the equilibrium is at most 1.102 times the cost of the global optimum. Note that when $m_I(0) = 0$ or $m_S(0) = 0$, the ratio is one because both the equilibrium and the socially optimal strategy consist in non-vaccinating.

We report the efficiency of the equilibrium for the mean field model in Figure 9. To obtain this figure, we consider all initial conditions $m_I(0), m_S(0) \in \{0, 1/100, \dots, 1\}$ and report the ratio between the cost of the mean field equilibrium and the cost of the socially optimal strategy. We observe that, as for the finite player model, the ratio is often close to 1. It is maximal when $m_S(0) \approx 0.67$




Figure 9: Ratio between the cost of the mean field equilibrium and the optimum mean field social cost. We plot the ratio as a function of the initial state $(m_S(0), m_I(0))$.

and $m_I(0) = 1/100$ (i.e. close to 0 but not equal to 0).

8 Conclusion

In this paper, we studied a mean field game model of a SIR dynamics where players can vaccinate. We showed that this game has a unique equilibrium that is of bang-bang type. We compared this equilibrium with the socially optimal strategy. We conclude that, if people decide by themselves when to vaccinate, then, in order to encourage a socially optimal behavior, vaccination should be subsidized.

In addition of its *per se* interest, we think that this problem is a good example where a mean field equilibrium can be computed under almost closed form while an analysis of the N -player Nash equilibrium seems impossible.

Acknowledgments

Josu Doncel has received funding from the Department of Education of the Basque Government through the Consolidated Research Group MATHMODE (IT1294-19), from the Marie Skłodowska-Curie grant agreement No 777778 and from the Spanish Ministry of Science and Innovation project with reference PID2019-108111RB-I00 (FEDER/AEI).

References

- [1] R. M. Anderson, R. M. May, and B. Anderson. *Infectious diseases of humans: dynamics and control*, volume 28. Wiley Online Library, 1992.
- [2] C. T. Bauch and D. J. Earn. Vaccination and the theory of games. *Proceedings of the National Academy of Sciences of the United States of America*, 101(36):13391–13394, 2004.
- [3] E. Bayraktar and A. Cohen. Analysis of a finite state many player game using its master equation. *arXiv preprint arXiv:1707.02648*, 2017.
- [4] H. Behncke. Optimal control of deterministic epidemics. *Optimal control applications and methods*, 21(6):269–285, 2000.
- [5] D. P. Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 1995.
- [6] R. Carmona, F. Delarue, and A. Lachapelle. Control of McKean–Vlasov dynamics versus mean field games. *Mathematics and Financial Economics*, 7(2):131–166, 2013.
- [7] A. Cecchin and M. Fischer. Probabilistic approach to finite state mean field games. *Applied Mathematics & Optimization*, Mar 2018.
- [8] A. Cecchin and G. Pelino. Convergence, fluctuations and large deviations for finite state mean field games via the master equation. *Stochastic Processes and their Applications*, 129(11):4510–4555, 2019.
- [9] O. Diekmann and J. A. P. Heesterbeek. *Mathematical epidemiology of infectious diseases: model building, analysis and interpretation*, volume 5. John Wiley & Sons, 2000.
- [10] J. Doncel, N. Gast, and B. Gaujal. Discrete Mean Field Games: Existence of Equilibria and Convergence. *Journal of Dynamics and Games*, 6(3):1–19, 2019.
- [11] P. J. Francis. Optimal tax/subsidy combinations for the flu season. *Journal of Economic Dynamics and Control*, 28(10):2037–2054, 2004.
- [12] N. Gast. Expected values estimated via mean-field approximation are $1/n$ -accurate. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 1(1):1–26, 2017.
- [13] N. Gast and B. V. Houdt. A Refined Mean Field Approximation. In *ACM SIGMETRICS 2018*, page 1, Irvine, United States, June 2018.
- [14] P.-Y. Geoffard and T. Philipson. Disease eradication: private versus public vaccination. *The American Econ. Review*, 87(1):222–230, 1997.

- [15] D. A. Gomes, J. Mohr, and R. R. Souza. Discrete time, finite state space mean field games. *Journal de Mathématiques Pures et Appliquées*, 93(3):308–328, 2010.
- [16] D. A. Gomes, J. Mohr, and R. R. Souza. Continuous time finite state mean field games. *Applied Mathematics & Optimization*, 68(1):99–143, 2013.
- [17] O. Guéant. Existence and uniqueness result for mean field games with congestion effect on graphs. *Applied Mathematics & Optimization*, 72(2):291–303, 2014.
- [18] E. Hubert and G. Turinici. Nash-mfg equilibrium in a sir model with time dependent newborn vaccination. 2016.
- [19] W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. In *Proceedings of the Royal Society of London A: mathematical, physical and engineering sciences*, volume 115, pages 700–721. The Royal Society, 1927.
- [20] V. N. Kolokoltsov, J. Li, and W. Yang. Mean field games and nonlinear markov processes. *arXiv preprint arXiv:1112.3744*, 2011.
- [21] 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.
- [22] L. Laguzet and G. Turinici. Individual vaccination as Nash equilibrium in a SIR model: the interplay between individual optimization and societal policies. Hal-01100579, version 1. Available: <https://hal.archives-ouvertes.fr/hal-01100579v1>, 2015.
- [23] R. Morton and K. H. Wickwire. On the optimal control of a deterministic epidemic. *Advances in Applied Probability*, 6(4):622–635, 1974.
- [24] L. S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953.
- [25] E. Todorov. Optimal control theory. *Bayesian brain: probabilistic approaches to neural coding*, pages 269–298, 2006.