

On Efficiency in Mean Field Differential Games

Maximilian Balandat and Claire J. Tomlin

Abstract— We investigate the efficiency of Nash equilibria of a class of Mean Field Games. We focus on the stationary case with entry and exit of players, and derive an expression for the social cost at a Nash equilibrium, based on value function and agent density. We propose a model for a Mean Field Congestion Game, in which the agents’ control cost depends (locally) on the agent density. We present numerical results that show that the Nash equilibria of these games are inefficient in general. Also, we point out an interesting paradox, which can be seen as a continuous analogue of Braess’s paradox known from selfish routing games. Finally, we cast the welfare maximization problem as a PDE-constrained optimization problem.

I. INTRODUCTION

Differential games model interactions of multiple non-cooperative agents in the context of a dynamical system. Computing Nash equilibria for general N -player games with non-trivial coupled costs quickly becomes intractable, as it amounts to finding solutions to a set of N coupled Hamilton-Jacobi-Bellman (HJB) PDEs [1]. To remedy this issue, there has recently been an increasing interest in Mean Field Games (MFG) [2], [3]. These games assume a homogenous population of agents who base their decisions only on their local state and the distribution of all other players (the mean field). It can be shown that in a suitable continuum limit (w.r.t. the number of players) the Nash equilibria of such a game are given by the solutions to a set of coupled PDEs [2].

Considering the important question of efficiency of equilibria, there has been relatively little work on this topic in the context of MFGs. Efficiency of equilibria in the particular case of Linear Quadratic Gaussian (LQG) games was addressed in [3]. In [4], variational techniques were employed to examine the efficiency of equilibria of a synchronization game between oscillators. Recently, [5] exploited a mean field structure to develop a decentralized implementation of the socially optimal control. This however assumed cooperative agents and is therefore quite different from our setup.

In this paper we investigate the efficiency of Nash equilibria of stationary MFGs in which players are constantly entering and exiting the state space. We formulate a Mean Field Congestion Game (MFCG), which models the interaction of agents whose “cost of movement” depends on the local agent density, similar to [6]. However, in [6] there was no discussion of efficiency. In the time-invariant, discounted cost setting, the solution to this game is a stationary one. Based on

the agents’ value function we derive a measure for the social cost, i.e. the overall cost for the population. We provide a numerical example, which shows that the outcome at a Nash equilibrium is generally inefficient. Moreover, we point out an interesting paradox that can be interpreted as a continuous analogue of Braess’s paradox known from selfish routing games [7]. Finally, we formulate the welfare optimization problem as a PDE-constrained optimization problem.

Outline: Section II recalls some basics on stationary MFGs. In section III we introduce the MFCG model, discuss some of its properties and provide a simple numerical example. Our main results on social cost and (in)efficiency of Nash equilibria are presented in section IV. Section V formulates the welfare maximization problem.

II. MEAN FIELD GAMES

As shown in [2], it is possible to derive a MFG as the limit ($N \rightarrow \infty$) of an anonymous N -player differential game with homogeneous players. In the continuum limit, the agent population is described by a density function over the state space. For the purpose of our discussion, however, we will start directly from a continuum model. Although there is really no notion of “a single” player, it is still helpful to retain this mental picture to get an intuitive understanding (as with the Eulerian and Lagrangian viewpoints in fluid dynamics).

A. Domain, agent dynamics and cost functions

In most of the the mathematical literature [2], [8], [9], [10], [11], MFGs are analyzed on simple domains under very simple boundary conditions. A frequent assumption is that the game takes place on a periodic state space. While reasonable from a mathematical perspective, this assumption is not too interesting from an engineering point of view.

We consider the state space $\mathcal{X} \subset \mathbb{R}^n$ to be an open, connected set with boundary $\partial\mathcal{X} \in C^1$ a.e.. Let $\mathcal{B} \subset \partial\mathcal{X}$ (“birth”) and $\mathcal{D} \subset \partial\mathcal{X}$ (“death”) with $\mathcal{B} \cap \mathcal{D} = \emptyset$ denote entry and exit of the state space, respectively. The agents’ desire to reach the exit will be encoded in their cost functions. In [6], the authors also consider a bounded domain, but use simpler boundary conditions and model entry and exit of players through sources and sinks inside the domain.

Let $m : \mathbb{R}_+ \times \mathcal{X} \rightarrow \mathbb{R}_+$ denote the agent density of the population over the state space, i.e. $m(t, x) =: m_t(x)$ is the agent density at point x at time t . We assume the dynamics of an agent are described by the Itô SDE

$$dX_t = \alpha_t dt + \Sigma dW_t + dN_t(X_t) \quad (1)$$

where $X_t \in \mathcal{X}$ and $\alpha_t \in \mathcal{A}$ are the agent’s state and control at time t , respectively, with $\mathcal{A} \subset \mathbb{R}^n$ the set of admissible

M. Balandat and C. Tomlin are both with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, USA [balandat, tomlin]@eecs.berkeley.edu

This work has been supported in part by NSF under CPS:ActionWebs (CNS-931843), by ONR under the HUNT (N0014-08-0696) and SMARTS (N00014-09-1-1051) MURIs and by grant N00014-12-1-0609, by AFOSR under the CHASE MURI (FA9550-10-1-0567).

control inputs. W_t is a standard n -dimensional Brownian motion and $\Sigma \in \mathbb{R}^{n \times n}$ is a diffusion coefficient matrix such that $\Sigma \Sigma^T \succ 0$. The reflection term $N_t(X_t)$ ensures that the process X_t does not leave \mathcal{X} through $\partial\mathcal{X} \setminus \{\mathcal{B} \cup \mathcal{D}\}$.

Remark 1: The dynamics (1) we assume are very simple, as agents can directly control their drift. It would be interesting to allow for more general dynamics, but this would require a careful treatment of the issue of controllability.

B. The stationary Mean Field Game equations

Define $\tau = \inf \{t \geq t_0 \mid X_t \in \mathcal{D}\}$ to be the (random) exit time of an agent whose state at time t_0 is $x_0 \in \mathcal{X}$. Denote by $\alpha := \{\alpha_t : t \geq t_0\}$ the agent's control profile. Suppose the agent aims to minimize the cost functional

$$J(x_0; \alpha) = \mathbf{E} \left[\int_{t_0}^{\tau} l(X_t, \alpha_t, m_t) e^{-\gamma(t-t_0)} dt + g(X_\tau) \mid X_{t_0} = x_0 \right] \quad (2)$$

with $l : \mathcal{X} \times \mathcal{A} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ the exit cost function, $g : \mathcal{D} \rightarrow \mathbb{R}$ the terminal cost function and $\gamma > 0$ the discount factor. We have $m_t = m_t(X_t)$, i.e. m_t is the density the agent experiences along its trajectory, and $\alpha_t = \alpha_t(X_t, m_t)$.

Remark 2: The dependence of the cost on m in (2) is highly localized. One may also want to consider non-local dependencies of the form $l(X_t, \alpha_t, m_t) = L_{[m_t]}(X_t, \alpha_t)$, where $L : \Delta(\mathcal{X}) \rightarrow C(\mathcal{X}, \mathcal{A})$, which for example could model local averaging. For our purposes of modeling congestion, a localized dependence seems plausible.

If $m = m(t, x)$ were fixed, then minimizing (2) over α would be a classic stochastic optimal control problem [12]. Since in this work we are interested in stationary mean field equilibria, suppose that $m_t(x) = m(x)$. Then one can show (see [8] for a rigorous derivation) using Itô's calculus and dynamic programming that the value function $v(x_0) = \inf_{\alpha \in \mathcal{A}} J(x_0; \alpha)$ satisfies the following stationary Hamilton-Jacobi-Bellman (HJB) PDE:

$$\frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 v) + H(x, Dv, m) - \gamma v = 0 \quad (3)$$

with Hamiltonian $H(x, p, m) = \inf_{a \in \mathcal{A}} \{p^T a + l(x, a, m)\}$ and boundary condition $v(x) = g(x)$, $\forall x \in \mathcal{D}$. An agent's optimal control input as function of its state and local density is $\alpha^*(x, m) = \arg \min_{a \in \mathcal{A}} \{(Dv)^T a + l(x, a, m)\}$.

Contrary to our previous thought experiment, the density in (3) is of course not fixed, since the agents themselves constitute the population. Hence it is their very choice of α that will determine the density m . A Nash equilibrium of the game is a pair (α, m) where the two are consistent.

A Nash eq. of a MFG is sometimes defined as the limit of a Nash eq. of a symmetric and anonymous N -player game. In this paper, we define it directly for the continuum formulation. Consider an agent population using the stationary control profile $\alpha : \mathcal{X} \rightarrow \mathcal{A}$ and denote by m^α the (stationary) density resulting from this choice. Define

$$J(x_0; \hat{\alpha}, \alpha) = \mathbf{E} \left[\int_{t_0}^{\tau} l(X_t, \hat{\alpha}_t, m_t^\alpha) e^{-\gamma(t-t_0)} dt + g(X_\tau) \mid X_{t_0} = x_0 \right]$$

The function $J(x_0; \hat{\alpha}, \alpha)$ can be interpreted as the expected overall cost for an "infinitely small" agent who at time t_0 is located at x_0 and uses the control profile $\hat{\alpha}$.

Definition 1: A stationary control profile α is called a *symmetric stationary mean field Nash equilibrium* if

$$J(x; \tilde{\alpha}, \alpha) \leq J(x; \alpha, \alpha) \quad (4)$$

for all $x \in \mathcal{X}$ and all $\tilde{\alpha} \in \mathcal{A}$.

In the following we refer to a symmetric stationary mean field Nash equilibrium simply as a Nash equilibrium.

The evolution of a density function under a controlled diffusion process such as X_t in general is described by a Fokker-Planck-Kolmogorov (FPK) PDE, which under appropriate boundary conditions admits a stationary solution. In our case, m satisfies the stationary FPK equation

$$-\frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 m) + \text{div}(m \alpha^*) = 0 \quad (5)$$

Introducing $F(x) := -\frac{1}{2} \Sigma \Sigma^T D m + m \alpha^*$ as the agent flow, we can rewrite (5) as $\text{div} F = 0$.

Equations (3) and (5) are coupled: The equilibrium density, satisfying (5), appears in (3) through the Hamiltonian H . Conversely, the value function (via its gradient) appears through the agents' optimal control in (5). It can be shown [2] that a Nash equilibrium of the game is described by a fixed point (v, m) that jointly solves (3) and (5).

C. Boundary conditions

The system of coupled PDEs (3) and (5) describes the equilibrium solution inside \mathcal{X} . To fully specify the problem we need to impose boundary conditions:

1) We require that agents can enter or leave the state space only through \mathcal{D} and \mathcal{B} , respectively. This translates to

$$\frac{\partial m}{\partial \vec{n}} = Dm(x) \cdot \vec{n}(x) = 0, \quad \forall x \in \partial\mathcal{X} \setminus \{\mathcal{B} \cup \mathcal{D}\} \quad (6)$$

where $\vec{n}(x)$ is the outward pointing surface normal of $\partial\mathcal{X}$. Note that because $\text{div} F = 0$ for all $x \in \mathcal{X}$, it follows from the divergence theorem and the fact that $\mathcal{B} \cap \mathcal{D} = \emptyset$ that

$$\int_{\mathcal{D}} F(x) \cdot \vec{n}(x) + \int_{\mathcal{B}} F(x) \cdot \vec{n}(x) = 0$$

The above is just a conservation condition stating that in the stationary case the net flow of agents into \mathcal{X} is zero.

2) We are interested in the Nash equilibria of the MFCG for a fixed overall agent flow f_{in} through the state space. This can be captured by the following boundary condition:

$$\int_{\mathcal{B}} F(x) \cdot \vec{n}(x) = -f_{in} \quad (7)$$

3) Finally, we assume that there is a certain cost associated with leaving the state space. This translates to the following boundary condition on the value function:

$$v(x) = g(x), \quad \forall x \in \mathcal{D} \quad (8)$$

D. Some comments on the finite horizon case

Setting $\gamma = 0$, $t_0 = 0$ and replacing τ by some fixed $T > 0$ in (2) yields a finite horizon MFG. Value function, density and agent flow then depend explicitly on time, and additional terms $\partial_t v$ and $\partial_t m$ appear in (3) and (5). The boundary conditions (6) and (7) become time-dependent and (8) is replaced by $v(T, x) = \tilde{g}(x), \forall x \in \mathcal{X}$. The function $m(t, x)$ describes the evolution of the agent density over time.

For brevity we will focus on the stationary case in the following. However, under appropriate modifications, our ideas extend also to the finite horizon case (see also Remark 3).

III. A MEAN FIELD CONGESTION GAME FRAMEWORK

In this section we introduce a simple model for a Mean Field Congestion Game (MFCG). In this model, the agents' "cost of movement" is assumed to depend (locally) on the agent density. This is motivated by a number of different applications: In traffic, for example, it is reasonable to assume a higher cost of moving when the local traffic density is high. In this case agents are inflicting a negative externality on their nearby peers. Conversely, one may want to consider positive externalities, which are of interest in some economic models. Here, we focus on the congestion case.

A. Cost function and MFCG equations

Consider a running cost of the form

$$l(x, \alpha, m) = \frac{1}{2} \alpha^T R(x, m) \alpha + k(x) \quad (9)$$

The first term in (9) describes the agents' cost of movement. The matrix-valued function $R : \mathcal{X} \times \mathbb{R}_+ \rightarrow \mathcal{S}_{++}^n$ (with \mathcal{S}_{++}^n the p.d. cone in \mathbb{R}^n) is monotonic w.r.t. the local density, i.e.

$$R(x, m_1) \succeq R(x, m_2), \quad \forall m_1 \geq m_2, \quad (10)$$

modeling congestion. The second component $k : \mathcal{X} \rightarrow \mathbb{R}_+$ of l encodes the agents' disutility for not having reached their goal yet. We call k the "pressure function", as it describes the pressure for agents to move towards the exit.

The above model, while similar to the one in [6], also allows for explicit dependence of the cost of movement on the state x and, through R , on the direction of α . That is, $|\alpha_1| = |\alpha_2| \not\Rightarrow l(x, \alpha_1, m) = l(x, \alpha_2, m)$.

With (9) one obtains the following set of MFCG equations:

$$-\frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 v) + \frac{1}{2} (Dv)^T R(x, m)^{-1} Dv + \gamma v = k(x) \quad (11)$$

$$\frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 m) + \text{div}(m R(x, m)^{-1} Dv) = 0 \quad (12)$$

B. Numerical solution

Numerical methods for MFG are a relatively new research topic [10] and no standard solvers are available for the coupled system (11) and (12) with boundary conditions as in section II-C. We implemented an algorithm that is based on fully discretizing both PDEs using a finite-difference scheme, and solving the resulting system of nonlinear algebraic equations using a trust-region solver. While this approach is straightforward, it clearly suffers from the curse of dimensionality.

At this point, we do not have theoretical results on the convergence of the algorithm. However, we have found the method to work well at least on simple numerical examples.

C. MFCGs in one dimension

Consider the special case when $\mathcal{X} = (0, 1)$ with entry $\mathcal{B} = \{0\}$ and exit $\mathcal{D} = \{1\}$. This very simple model could for example be used to describe traffic on a stretch of highway. In the finite horizon time-varying case, a similar model has been used in an economic context to analyze the dynamic equilibrium in a "choice of technologies" game [13]. As we will see, even this simple one-dimensional stationary model exhibits interesting and sometimes surprising properties w.r.t. the efficiency of its equilibria.

The condition $\text{div } F(x) = 0$ for $x \in \mathbb{R}$ simply means that the flow is constant over \mathcal{X} . In particular, this constant has to equal in- and outflow, which reduces (12) to $F(x) = f_{in}$. The MFCG equations thus read

$$-\frac{\sigma^2}{2} v'' + \frac{1}{2} \frac{(v')^2}{r(x, m(x))} + \gamma v = k(x) \quad (13)$$

$$\frac{\sigma^2}{2} m' + \frac{m v'}{r(x, m(x))} = -f_{in} \quad (14)$$

The domain of g reduces to the single point $\{1\}$, so $g = v(1)$ describes the cost incurred when an agent leaves. W.l.o.g. we assume $v(1) = 0$. Notice that the one-dimensional MFCG equations are a set of ODEs rather than PDEs.

We now consider the one-dimensional deterministic limit case, i.e. when $\sigma \rightarrow 0$. Under certain regularity conditions on the agent cost functions it is possible to compute a Nash eq. of the stationary MFCG by solving a single ODE with implicitly defined boundary condition.

Proposition 1: Consider a 1-dim. MFCG with $\sigma = 0$ and suppose k' , $\partial_x r$ and $\partial_m r$ exist $\forall x, m$. Suppose further that

$$1) \text{ the function } h_1(m) = r(1, m) f_{in}^2 - 2k(1) m^2 \quad (15)$$

has a unique positive root m_1 .

2) the function

$$h_2(x, m) = \left[\frac{\frac{2m}{f_{in}} (\gamma r(x, m) + \frac{m}{f_{in}} k'(x)) - \partial_x r}{m \partial_m r(x, m) - 2r(x, m)} \right] m \quad (16)$$

is p/w cont. in x and Lipschitz in m , $\forall x \in \mathcal{X}, m > 0$.

Then the stationary one-dim. MFCG has a unique Nash equilibrium. Further, the eq. density is given by the solution of $m'(x) = h_2(x, m)$ with boundary condition $m(1) = m_1$.

Proof: Letting $\sigma = 0$ and plugging (13) into (14) yields

$$v = \gamma^{-1} [k(x) - 0.5 r(x, m(x)) f_{in}^2 m^{-2}] \quad (17)$$

Using that $v(1) = 0$ we obtain (15) by rearranging (17). For a (unique) solution to exist we need h_1 to have a (unique) positive root. Differentiating (17), using (14) with $\sigma = 0$ we obtain (16) after some algebra. Uniqueness of the solution of (16) is established via the Fundamental Theorem of Differential Equations (since m_1 is unique). ■

For small σ , the solution of the deterministic limit case provides a good initial guess, which can be used to warm-start the solver based on discretizing the equations.

D. A simple numerical example

We present a simple numerical example in order to illustrate the Nash equilibria of the game for different agent cost functions. We fix the parameters $\gamma = 0.15$, $\sigma = 0.2$ and $f_{in} = 0.5$. The agent cost functions r and k are given by:

	$r(x, m)$	$k(x)$
—	$1 + 2m^2$	$0.75 + 0.75x$
- - -	$1 + 2m^2$	$1.75 - 0.75x$
- · - ·	$1 + 4m^2 + 0.75 \sin(2\pi x)$	1.5

The Nash eq. solutions for the different cost functions are shown in Fig. 1 (left: Eulerian, right: Lagrangian). The cost of movement (characterized by r) in the games corresponding to the solid and dashed lines depends only on $m(x)$ and not explicitly on x . If the pressure function k increases with x (solid), the eq. density decreases with x , and vice versa (dashed), as we would expect. In the game corresponding to the dash-dotted lines, the cost of movement has sinusoidal shape. We see that α^* is large when the cost of movement is low and vice versa, which complies with our intuition. However, α^* is not sinusoidal since r is non-linear in the density. Note also the difference in expected travel time $\mathbf{E}[\tau | x_0=0]$ and how the curvature of the value function v is related to slope of the pressure function k .

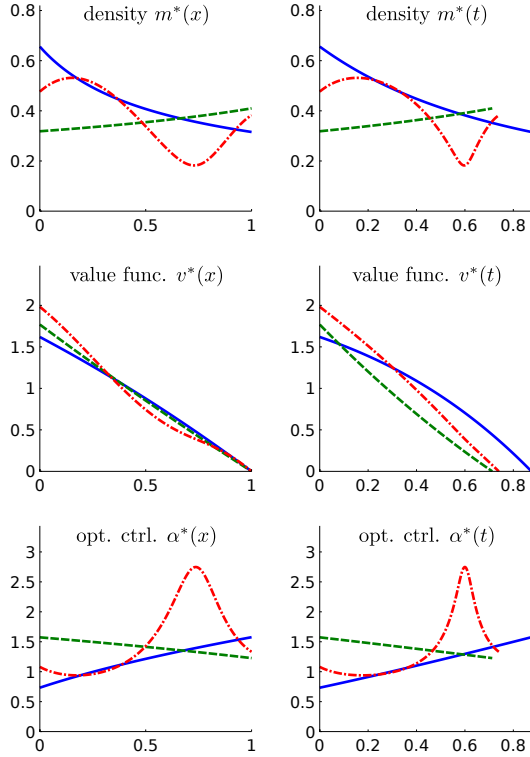


Fig. 1. Nash equilibrium solutions for different cost functions

IV. SOCIAL COST AND EFFICIENCY IN STATIONARY MFG

Central to the question of efficiency is the notion of the *social cost* – the overall cost for society. While in many games it is quite obvious what the social cost is, this is not the case in our model as (i) we are dealing with a continuum of agents, (ii) agents are entering and exiting the state space and (iii) stationarity requires a normalization w.r.t. time.

A. Social cost in stationary MFCG

We propose to use the expected cost of the overall population per time unit, call it C , as the measure for the social cost of a stationary Mean Field equilibrium.

Theorem 1: Consider a Nash eq. of a stationary MFG described by (v, m) solving (3) and (5). Then the expected cost incurred by the population per time unit is given by

$$C = - \int_{\mathcal{X}} [(Dv)^T \alpha + \frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 v)] dm \quad (18)$$

Proof: The expected cost \tilde{C} invoked by a single agent at point x_0 in some time interval Δt is given by $\tilde{C}(x_0, \Delta t) = \mathbf{E}[v(X_t) - v(X_{t+\Delta t}) | X_t = x_0]$. Due to stationarity this cost is independent of t . Normalizing by Δt and taking the limit $\Delta t \rightarrow 0$ yields $c(x_0) := \lim_{\Delta t \rightarrow 0} \tilde{c}(x_0, \Delta t)$. Applying Itô's lemma and using the dynamics (1) we obtain (formally)

$$\begin{aligned} dv(X_t) &= (Dv)^T dX_t + \frac{1}{2} dX_t^T D^2 v dX_t \\ &= ((Dv)^T \alpha_t + \frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 v)) dt + (Dv)^T \Sigma dW_t + O(dt^{\frac{3}{2}}) \end{aligned}$$

where we have used that $(dW_t)_i (dW_t)_j = \delta_{ij} dt$ (in particular: $dW_t dW_t^T = Idt$). Hence we have

$$\begin{aligned} v(X_{t+\Delta t}) &= v(X_t) + \int_t^{t+\Delta t} ((Dv)^T \alpha_s + \frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 v)) ds \\ &\quad + \int_t^{t+\Delta t} (Dv)^T \Sigma dW_s + O(\Delta t^{\frac{3}{2}}) \end{aligned}$$

Rearranging, taking conditional expectations and using the fact that the expectation of the Itô integral is zero we get

$$\begin{aligned} \mathbf{E}[v(X_t) - v(X_{t+\Delta t}) | X_t = x_0] &= \mathbf{E}\left[- \int_t^{t+\Delta t} ((Dv)^T \alpha_s + \frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 v)) ds + O(\Delta t^{\frac{3}{2}}) \mid X_t = x_0\right] \end{aligned}$$

for all $\Delta t \geq 0$. In particular, this implies that

$$c(x) = -(Dv(x))^T \alpha - \frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 v(x)) \quad (19)$$

The social cost C is the instantaneous cost invoked by the overall population, hence (18) is obtained by integrating $c(x)$ with respect to the equilibrium density m . ■

We want to emphasize that Theorem 1 makes no assumptions about the particular model used. Rather, it applies to any stationary game whose solution is described by a stationary value function and control. In particular, if α in (18) is replaced by the drift resulting from the optimal control, it applies also to games with general dynamics.

Corollary 1: If $n=1$ the cost (18) can be written as

$$C = v(0) f_{in} - \frac{\sigma^2}{2} [m(1) v'(1) - m(0) v'(0)] \quad (20)$$

Proof: Conservation of flow for $n=1$ is described by $-\frac{\sigma^2}{2} m' + m\alpha = f_{in}$. Using this in (18) yields

$$\begin{aligned} C &= - \int_{\mathcal{X}} \left[v' f_{in} + \frac{\sigma^2}{2} (m' v' + m v'') \right] dx \\ &= f_{in} [v(x)]_1^0 - \frac{\sigma^2}{2} \int_0^1 (m' v' + m v'') dx \end{aligned}$$

Performing integration by parts on the term $\int m' v' dx$ and using the BC $v(1) = 0$ we obtain (20). ■

The expression (20) is more intuitive than the general form (18): In the deterministic limit $\sigma = 0$ the social cost is the expected cost for a newly arriving agent times the arrival rate f_{in} . In the stochastic case, there is an additional term accounting for the net difference in diffusion-induced flow.

Remark 3: For the finite horizon MFG from section II-D the social cost \tilde{C} can be expressed as

$$\tilde{C} = \int_{\mathcal{X}} v(0, x) dm_0 - \int_0^T \int_{\mathcal{B}} v(t, x) F(t, x) \cdot \bar{n}(x) dx dt \quad (21)$$

where the first term describes the cost invoked by agents present in \mathcal{X} at $t = 0$ (with density $m_0 = m(0, x)$), and the second term captures the cost invoked by agents entering the state space through \mathcal{B} during $(0, T)$.

B. Inefficiency of Nash Equilibria

In this section we provide a numerical example which shows that Nash eq. of MFCG are generally inefficient. We consider a one-dim. stationary MFCG as in section III-C with $f_{in} = 0.5$, $\sigma = 0.2$, $\gamma = 0.1$, $k(x) = 1 - (1 - x)/4$ and $r(x, m) = 1 + 3m^2$. Denote by α^* and C^* control profile and social cost at Nash eq., respectively. We investigate the social cost C if the population (as a whole) were to deviate from α^* . Specifically, consider $\alpha(x) = \alpha^*(x) + \eta_1 + \eta_2 \sin(2\pi x)$ where $\eta = (\eta_1, \eta_2) \in \mathbb{R}^2$, i.e. a perturbation of the Nash eq. profile α^* in two directions (in functional space). We refer to section (V) on how to compute the social cost for such an exogenously chosen control α .

Fig. 2 shows the normalized social cost $C(\eta)/C^*$. With respect to the chosen perturbations, the cost has a local minimum near $\hat{\eta} = (0.65, 0.1)$ with $C(\hat{\eta})/C^* = 0.89$. This shows that the control profile α^* is not socially optimal. The interpretation here is that because η_1 is quite large, all agents are moving faster through the state space. While this renders agents close to the exit at $x = 1$ worse off (they incur a higher cost of movement but do not gain much, as they would leave the state space soon anyway), agents further left are better off. Overall, society has a net benefit.




Fig. 2. Normalized social cost as a function of η

C. A continuous analogue of Braess's paradox

Braess's paradox is a well known paradox in selfish routing games [7]. It shows that removing links from a network in which agents choose paths selfishly so as to minimize their cost can lead to a decrease in social cost (in fact decrease the cost for every agent). One can observe

an analogous situation in MFCG: Strictly increasing the cost functions for every agent may decrease social cost.

We again consider a 1D MFCG as in section III-C with $f_{in} = 0.5$, $\sigma = 0.15$, $\gamma = 0.2$, $k(x) = 1 + (1 - x)$ and $r(x, m) = 1 + 5m^2$. We modify the agents' pressure function as follows: $k_{mod}(x) = k(x) + \beta x^3$. Note that for $\beta > 0$ we have $k_{mod}(x) > k(x)$, i.e. the agents' disutility from not having reached the exit is strictly increased.

Denote by $C^*(\beta)$ the social cost at the Nash eq. of the modified game. Numerically we find that $C^*(\beta) < C^*(0)$ for $\beta \in (0, 1.175)$, meaning that the social cost at the Nash equilibrium of the modified games is decreased, even if the agents' cost function has been strictly increased.

Some explanation for this counterintuitive observation can be obtained from Fig. 3, which shows the Nash eq. density $m^*(x)$ and the associated optimal control profile $\alpha^*(x)$ for different β . For small β , the density close to $x = 1$ is quite high. Agents in this area have little incentive to exit the state space quickly, as $k(x)$ is rather small for $x \approx 1$. Hence the control α^* in this region is also small, resulting in a higher local density. As a result, agents further to the left will be incurring a higher cost when traveling through this congested area. In economic terms, the agents close to $x = 1$ impose a negative externality on those to their left.

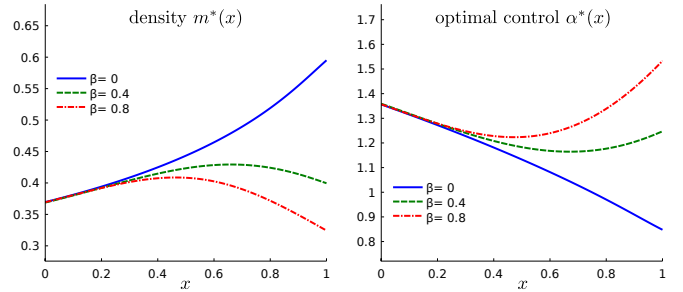


Fig. 3. Density and optimal control at Nash Eq. for different values of β

An increasing β corresponds to a growing desire of agents near $x = 1$ to leave. These agents will choose to move faster, thereby reducing the density (see Fig. 3) and reducing “blocking” of other agents. While a small part of the population close to $x = 1$ will have a higher cost for their *remaining* route, the majority of the population sees a decrease in cost, so that society as a whole has a net benefit. Note that the word *remaining* here is essential: At the new Nash eq. every agent entering at $x = 0$ will be better off.

Numerically one can find that $\beta = 0.4$ corresponds to a local minimum of $C^*(\beta)$. For values of β higher than 0.4 additional gains are consumed by the increase in cost, so that for $\beta > 1.175$ society is actually worse off.

With around 1.5%, the decrease of the social cost in the above example is rather small. However, this was a specific example with a specific modification to k , so one may expect this decrease to be more pronounced in other cases.

V. THE SOCIAL PLANNER PROBLEM

We would like to answer the following question: “How should a benevolent dictator choose the population's control

profile in order to minimize social cost?”. In welfare economics, this problem is known as the social planner problem.

A. The welfare maximization problem

Suppose the population was using some exogenously chosen stationary control profile $\alpha(x)$. Define the expected cost-to-go function \bar{v} for this fixed policy as:

$$\bar{v}(x_0; \alpha) = \mathbf{E} \left[\int_{t_0}^{\tau} e^{-\gamma(t-t_0)} l(X_t, \alpha_t, m_t) dt + g(X_\tau) \mid X_0 = x \right] \quad (22)$$

Proposition 2: The cost-to-go function \bar{v} given by (22) satisfies the following PDE:

$$\frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 \bar{v}) + (D\bar{v})^T \alpha + l(x, \alpha, m) - \gamma \bar{v} = 0 \quad (23)$$

with boundary condition $v(x) = g(x)$, $\forall x \in D$.

Proof: The PDE (23) is obtained by mimicking the dynamic programming derivation of the HJB equation. ■

Note that since the proof of Theorem 1 does not depend on (v, m) being a Nash eq, we can still use (18) with v replaced by \bar{v} to compute the social cost for a particular α . The social planner problem can then be formulated as the following PDE-constrained optimization problem:

$$\begin{aligned} \hat{C} = & \min_{\bar{v}, m, \alpha} - \int_{\mathcal{X}} \left[(D\bar{v})^T \alpha + \frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 \bar{v}) \right] m dx \\ \text{s.t. } & \frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 \bar{v}) + (D\bar{v})^T \alpha + l(x, \alpha, m) - \gamma \bar{v} = 0 \\ & -\frac{1}{2} \text{tr}(\Sigma \Sigma^T D^2 m) + \text{div}(m \alpha) = 0 \end{aligned}$$

The above problem is hard: The variable α is infinite dimensional, and we do not have a priori information about convexity. Still, there do exist tools to tackle this kind of problem; we are focusing on solving some simple cases using adjoint-based optimization techniques [14].

B. Extensions and future research directions

The social planner problem is important, as it provides a lower bound on the achievable social cost. It is central to the Price of Anarchy [7], [15], which characterizes how “bad” a Nash eq. of a game may be compared to the social optimum. An interesting question then is whether it is possible to obtain bounds on the Price of Anarchy in terms of parameters of the game, similar in spirit to those derived in [7] for routing games. The problem here is of course much more complex.

Another interesting avenue for future research is the problem of designing incentives through pricing or subsidizing certain behavior, so that the Nash eq. of the modified game become more desirable from a social point of view. Similar to the social planner problem, this can also be formulated as a PDE-constrained optimization problem in which the social cost is optimized over modifications (taxes or subsidies) to the agents’ cost functions subject to PDE constraints describing the Nash equilibria of the modified game.

VI. CONCLUSION

Building on the Mean Field Games framework, we proposed a model for stationary Mean Field Congestion Games with entry and exit of players. We implemented a numerical algorithm for solving these games based on the discretization of the corresponding system of coupled PDEs. Our two main contributions were to derive an expression for the social cost of stationary MFG and to show that the Nash equilibria of these games are generally inefficient. We also identified an interesting continuous analogue of Braess’s paradox. Finally, we formulated the social planner problem as a PDE-constrained optimization problem. We believe that this work is one of the first to investigate the efficiency of mean field equilibria in a more general setting, and that it raises interesting further questions on issues like the Price of Anarchy and the design of incentives for Mean Field Games.

ACKNOWLEDGMENTS

Thanks to Lillian Ratliff and Pushkin Kachroo for valuable discussions about parts of this work.

REFERENCES

- [1] A. Bensoussan and J. Frehse, “Nonlinear elliptic systems in stochastic game theory,” *Journal für die reine und angewandte Mathematik (Crelles Journal)*, vol. 350, pp. 23–67, 1984.
- [2] J.-M. Lasry and P.-L. Lions, “Mean field games,” *Japanese Journal of Mathematics*, vol. 2, pp. 229–260, 2007.
- [3] M. Huang, P. Caines, and R. Malhamé, “Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ϵ -nash equilibria,” *Automatic Control, IEEE Transactions on*, vol. 52, no. 9, pp. 1560–1571, Sep 2007.
- [4] H. Yin, P. Mehta, S. Meyn, and U. Shanbhag, “On the efficiency of equilibria in mean-field oscillator games,” in *American Control Conference (ACC)*, Jul 2011, pp. 5354–5359.
- [5] M. Huang, P. Caines, and R. P. Malhamé, “Social optima in mean field lqg control: Centralized and decentralized strategies,” *Automatic Control, IEEE Transactions on*, vol. 57, no. 7, pp. 1736–1751, 2012.
- [6] A. Lachapelle and M.-T. Wolfram, “On a mean field game approach modeling congestion and aversion in pedestrian crowds,” *Transportation Research Part B: Methodological*, vol. 45, no. 10, pp. 1572–1589, 2011.
- [7] T. Roughgarden and E. Tardos, “How bad is selfish routing?” *J. ACM*, vol. 49, no. 2, pp. 236–259, Mar 2002.
- [8] P.-L. Lions, “Théorie des jeux de champ moyen et applications,” Cours au Collège de France, 2008-2009.
- [9] O. Guéant, “A reference case for mean field games models,” *Journal de Mathématiques Pures et Appliquées*, vol. 92, no. 3, pp. 276–294, 2009.
- [10] Y. Achdou and I. Capuzzo-Dolcetta, “Mean field games: Numerical methods,” *SIAM Journal on Numerical Analysis*, vol. 48, no. 3, pp. 1136–1162, 2010.
- [11] O. Guéant, J.-M. Lasry, and P.-L. Lions, “Mean field games and applications,” in *Paris-Princeton Lectures on Mathematical Finance 2010*, ser. Lecture Notes in Mathematics. Springer, 2011, vol. 2003, pp. 205–266.
- [12] W. Fleming and H. Soner, *Controlled Markov processes and viscosity solutions*, 2nd ed., ser. Stochastic modelling and applied probability. Springer-Verlag, 2006.
- [13] A. Lachapelle, J. Salomon, and G. Turinchi, “Computation of mean field equilibria in economics,” *Math. Models Methods Appl. Sci.*, vol. 20, pp. 567–588, 2010.
- [14] A. Bensoussan, G. D. Prato, M. C. Delfour, and S. K. Mitter, *Representation and Control of Infinite Dimensional Systems (Systems & Control: Foundations & Applications)*, 2nd ed. Birkhäuser, 2007.
- [15] T. Başar and Q. Zhu, “Prices of anarchy, information, and cooperation in differential games,” *Dynamic Games and Applications*, vol. 1, pp. 50–73, 2011.