

Large Networks of Dynamic Agents: Consensus under Adversarial Disturbances

Dario Bauso
DICGIM, Università di Palermo
V.le delle Scienze, 90128 Palermo, Italy
E-mail: dario.bauso@unipa.it

Tamer Başar
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
1308 W. Main St., Urbana IL, USA
e-mail: basar1@illinois.edu

Abstract—This paper studies interactions among homogeneous social groups within the framework of large population games. Each group is represented by a network and the behavior described by a two-player repeated game. The contribution is three-fold. Beyond the idea of providing a novel two-level model with repeated games at a lower level and population games at a higher level, we also establish a mean field equilibrium and study state feedback best-response strategies as well as worst-case adversarial disturbances in that context.

I. INTRODUCTION

“Systems of systems” is a topic of primary interest in the European Union FP7 work program. Systems considered here are *social networks/groups*, which at an abstract level, are described by mathematical models that try to capture the main phenomena when a large number of individuals interact. The purpose of this current study is to analyze the mutual influences in a large number of homogeneous social networks; see Fig. 1. On each single network, the daily interactions among

up to the current time. The perspective we adopt is a worst-case one, in the sense that each network state is affected by an adversarial disturbance beyond the presence of a controlled input. Thus the resulting game on each network is a two-player repeated one, with vector payoffs. The first player selects the controlled input, the second player the adversarial disturbance. More specifically, given a current state distribution over the entire population (of networks), the controlled input of a single network tries to steer the state to the mean of the state distribution. This resembles herd behaviors or crowd-seeking attitudes in that certain social groups tend to mimic the behavior of other social groups. A similar behavior can be observed in financial markets under the name of “stock market bubbles”, which sees investors to emulate other investors. A third application area is in everyday decision making in that decisions are made on the basis of the observed information (past decisions), thus influencing successive decisions. Such a phenomenon is known as “cascaded information” [3], [8].

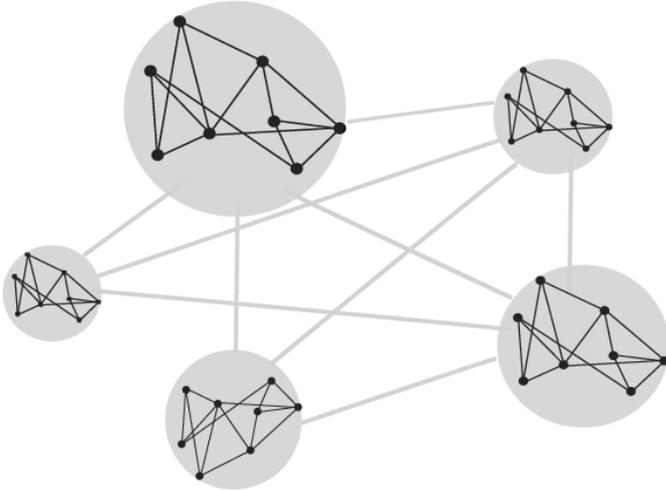


Fig. 1. Network of networks.

players is captured by a vector payoff, one component per each player. The state of the network is the cumulative payoff

Outline of the results. A first result in this paper involves providing a mean field game framework that captures the interactions among a large number of networks where herding behavior is rewarding. The novelty of the proposed model is in the two-layer structure. On a lower layer one observes two-player repeated games, one for each network. On a higher layer a large number of networks interact according to a mean field game model.

For the mean field game at hand, a second result is that we establish a *mean field equilibrium*, that is, a solution in terms of state feedback best-responses for the controlled inputs. In such solutions no player can benefit from a unilateral deviation, as in the case of Nash equilibrium strategies in differential games with a finite number of players [6]. At the equilibrium, controlled inputs exhibit a bang-bang structure when the associated state is far from the mean distribution (target value), and turn into linear controls in a neighborhood of the target value.

As a third contribution, microscopic and macroscopic analyses are carried out to show convergence properties of the population distribution. This is accomplished by resorting to Markov chain stability tools. Under certain assumptions, the population converges in mean and variance.

Related literature on mean field games. A first formulation of mean field games is due to J. M. Lasry and P. L. Lions [13]. The featuring aspect is that the strategies of a single agent are influenced by the mass of the other agents.

Economics, physics, biology, and network and production engineering are among the several domains where mean field games find applications (see [1], [7], [9], [10], [12], [15], [17], [21]).

The mathematical formulation of a mean field game consists of a system of two partial differential equations (PDEs). The first PDE is the *Hamilton-Jacobi-Bellman* (HJB) equation whose solution is the *value function* which is parametrized in the population distribution [7] (in the presence of disturbance the equation is better known as *Hamilton-Jacobi-Bellman-Isaacs* (HJBI)). The HJB equation is coupled with a second PDE, known as *Fokker-Planck-Kolmogorov* (FPK), defined on variable population distribution and parametrized in the value function [2], [13], [18], [20].

Explicit closed-form expressions for mean field equilibria are available only for a limited class of problems among which is the class of linear-quadratic mean field games, see [4]. As an alternative to explicit solutions, a variety of numerical approaches are available which hinge on discretization and finite difference approximations [2].

Evolutionary games represent another stream of literature strongly connected to mean field games and large population games ([16], [11], [19]). A first attempt to introduce dynamics is in [11] for a discrete time version of the game, which represents a precursor of mean field games.

More recently, and in the spirit of the present paper, the notion of robustness has been brought within the picture. Robust mean field games aim to achieve robust performance and/or stability in the presence of unknown disturbances when there is a large number of players; see [18], where relations with risk-sensitive games and risk-neutral games have been analyzed.

The rest of this paper is organized as follows. In Section II we formulate the problem and introduce the model. In Section III we establish a mean field equilibrium for the problem at hand. Finally, in Section IV we draw some conclusions.

II. THE MODEL

We consider a “large population” of *networks*, i.e., social groups characterized by a controlled time-varying behavior/state/characteristics. By “large population” we mean a large but finite number of *homogeneous networks*, namely, players’ interactions in each network are similar. Each network is subject to controlled inputs and adversarial disturbances. The large population game in its generic form appears as follows.

Given the following data, $f : \Delta(U) \times \Delta(W) \rightarrow [-1, 1]^n$, “bilinear”, U a finite control set and $\Delta(U)$ a set of mixed control actions, W a finite disturbance set and $\Delta(W)$ a set of mixed disturbance actions; $g : \mathbb{R}^n \times \mathbb{R}^n \rightarrow [0, +\infty[$, “regular”; $\Psi : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$, “regular”; \mathcal{U} and \mathcal{W} , sets of all measurable functions $u(\cdot)$ and $w(\cdot)$ from $[0, T]$ to $\Delta(U)$ and $\Delta(W)$,

respectively; $m : \mathbb{R}^n \times [0, +\infty[\rightarrow \mathbb{R}$, $(x, t) \mapsto m[x](t)$, probability density function, which satisfies $\int_{\mathbb{R}^n} m[x](t) dx = 1$ for every t ; each player solves the minimization problem:

$$\begin{aligned} & \inf_{u(\cdot) \in \mathcal{U}} \sup_{w(\cdot) \in \mathcal{W}} J(x, u(\cdot), w(\cdot), m\cdot) \\ & = \int_0^T g(X(t), \bar{m}(t)) dt + \Psi(X(T), m[\delta_{X(T)}](T)) \quad (1) \end{aligned}$$

$$dX(t) = f(u(t), w(t)) dt + \sigma dB(t), \quad X(0) = x, \quad (2)$$

where we denote by

- $X : [0, T] \rightarrow \mathbb{R}^n$, $t \mapsto X(t)$, the players’ state at time t , and x its initial state,
- $u : [0, T] \rightarrow \Delta(U)$, $t \mapsto u(t)$, the control at time t ;
- $w : [0, T] \rightarrow \Delta(W)$, $t \mapsto w(t)$, the disturbance at time t ;
- $\bar{m}(t) := \int_{\mathbb{R}^n} x m[x](t) dx$ the mean distribution at time t .

Note that the cost functional includes a mean field term to capture the mutual interaction between individuals and population. In this sense, observe that the distribution $m[x](t)$ represents the percentage of agents in state x at time t . Also, the cost functional J in (1), currently in a finite horizon form, can be replaced by infinite horizon discounted or average cost when appropriate. Also, the differential equation (2) in its stochastic form includes an infinitesimal vector Brownian motion $dB = \{dB_i\}^n$ (the i th component of dB is dB_i which is a scalar Brownian motion) weighted by a scalar coefficient σ , but turns into an ordinary differential equation if we take $\sigma = 0$.

We suppose that the distribution enters into the cost through its mean. Indeed, the running cost $g : \mathbb{R}^n \times \mathbb{R}^n \rightarrow [0, +\infty[$, $(x, \bar{m}) \mapsto g(x, \bar{m})$ is of the form:

$$g(x(t), \bar{m}(t)) = \frac{1}{2} (\bar{m}(t) - x(t))^T Q (\bar{m}(t) - x(t)).$$

The above cost describes the square deviation of an individual’s state from the mean state computed over the whole population. Such a cost reflects the willingness of the individuals to mimic the mean population behavior as it happens in herd behaviors or crowd-seeking attitudes.

To see how the model above describes players’ interactions over networks consider the following example borrowed from [14].

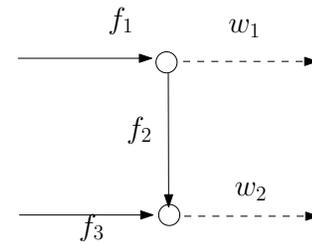


Fig. 2. Network system.

Example 1: Consider the system depicted in Fig. 2 representing two warehouses, three controlled flows and two uncontrolled flows. A unit of flow f_1 produces one unit of product X_1 per time unit. Similarly, flow f_2 uses one unit of

X_1 to produce one unit of X_2 per time unit. A unit of flow f_3 produces one unit of product X_2 per time unit. Uncontrolled flows w_1 and w_2 represent the exogenous demand of resources X_1 and X_2 , respectively. The associated dynamics then reads:

$$\begin{bmatrix} \dot{X}_1 \\ \dot{X}_2 \end{bmatrix} = f(u, w) = \underbrace{\begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & 1 \end{bmatrix}}_F \underbrace{\begin{bmatrix} f_1^t \\ f_2^t \\ f_3^t \end{bmatrix}}_{u(t)} - \underbrace{\begin{bmatrix} w_1^t \\ w_2^t \end{bmatrix}}_{w(t)}.$$

Now, suppose that flows can be processed only in batches and therefore take for instance $f_i \in \{-5, -2, 1, 6\}$, and $w_i \in \{-3, 2\}$.

Let us enumerate all the actions of players 1 and 2, so that we have $U = \{a_{11}, \dots, a_{1r}\}$ and $W = \{a_{21}, \dots, a_{2q}\}$ with $r = 4^3$ and $q = 2^2$, where a_{ij} denotes the j th action of player i .

The complete matrix of vector payoffs is then obtained from the following table, where each entry represents a possible vector payoff $f(u, w)$:

u/w	a_{21}	\dots	a_{2q}
a_{11}	$Fa_{11} - a_{21}$	\dots	$Fa_{11} - a_{2q}$
\vdots	\vdots		\vdots
a_{1r}			

As it will be clearer later, for our purposes we can simply extract from the above table the rows corresponding to the following four actions of player 1:

$$\begin{aligned} a_{11} &= (1, -2, 6), & a_{12} &= (1, -2, -5), \\ a_{13} &= (-5, 1, -5), & a_{14} &= (-5, 1, 6). \end{aligned}$$

For player 2, we consider the following four actions:

$$\begin{aligned} a_{21} &= (-3, -3), & a_{22} &= (2, -3), \\ a_{23} &= (-3, 2), & a_{24} &= (2, 2). \end{aligned}$$

Control and disturbance sets are then $U = \{a_{11}, \dots, a_{14}\}$ and $W = \{a_{21}, \dots, a_{24}\}$, respectively.

The following 4×4 matrix includes all possible vector payoffs $f(u, w)$, where $u \in U$ and $w \in W$ and $n = 2$:

$$\begin{pmatrix} (6, 7) & (1, 7) & (6, 2) & (1, 2) \\ (6, -4) & (1, -4) & (6, -9) & (1, -9) \\ (-3, -1) & (-8, -1) & (-3, -6) & (-8, -6) \\ (-3, 10) & (-8, 10) & (-3, 5) & (-8, 5) \end{pmatrix}. \quad (3)$$

Note that we can always normalize $f(\cdot)$ so as to have $(\Delta(U), \Delta(W)) \mapsto [-1, 1]^n$.

This then leads to the following mean field game system

$$\begin{cases} v_t(x, t) + \inf_{u \in \Delta(U)} \sup_{w \in \Delta(W)} \{f(u, w)v_x(x, t) + g(x, \bar{m}, t)\} \\ + \frac{\sigma^2}{2} v_{xx}^2(x, t) = 0 \text{ in } \mathbb{R}^n \times [0, T], \\ v(x, T) = \Psi(x) \quad \forall x \in \mathbb{R}^n \\ m_t(t) + \operatorname{div}(m \cdot f(u^*, w^*)) - \frac{\sigma^2}{2} m_{xx}^2 = 0, \quad m(0) = m_0 \\ u^*(t, x) = \operatorname{arginf}_{u \in \Delta(U)} \{f(u, w^*)v_x(x, t)\}. \\ w^*(t, x, u) = \operatorname{argsup}_{w \in \Delta(W)} \{f(u, w)v_x(x, t)\}. \end{cases}$$

The mean field game system appears in the form of two coupled PDEs intertwined in a forward-backward way. The first equation is the HJBI equation with variable $v(x, t)$ and parametrized in $m(\cdot)$. Given the boundary condition on final state (second equation), and assuming a given population behavior captured by $m(\cdot)$, the HJBI equation is solved backwards and returns the value function and best-response behavior of the individuals (fourth equation) as well as the worst adversarial response (fifth equation). The HJBI equation is coupled with a second PDE, known as *Fokker-Planck-Kolmogorov (FPK)* (third equation), defined on variable $m(\cdot)$ and parametrized in $v(x, t)$. Given the boundary condition on initial distribution $m(0) = m_0$, and assuming a given individual behavior described by u^* the FPK equation is solved forward and returns the population behavior time evolution $m(t)$.

Noting that the distribution enters into the cost through its mean, we can also simplify the above system through model reduction. Indeed, we can replace the FPK equation by an ordinary differential equation in the variable mean distribution and obtain the following system:

$$\begin{cases} v_t(x, t) + \inf_{u \in \Delta(U)} \sup_{w \in \Delta(W)} \{f(u, w)v_x(x, t) \\ + g(x, \bar{m}, t)\} + \frac{\sigma^2}{2} v_{xx}^2(x, t) = 0 \text{ in } \mathbb{R}^n \times [0, T], \\ v(x, T) = \Psi(x) \quad \forall x \in \mathbb{R}^n \\ \bar{m}'(t) = \int f(u^*, w^*)m(x, t)dx, \\ \bar{m}(0) = \bar{m}_0. \end{cases} \quad (4)$$

Any solution of the above system of equations is referred to as *mean field equilibrium*, as no individual benefits arise from deviating from the adopted strategy. A main question is the existence and uniqueness of equilibria for (4). The answer lies in most cases in numerical techniques and requires the implementation of a fixed point procedure. This entails assuming that $\bar{m} : [0, T] \rightarrow \mathbb{R}$ is continuous and $\bar{m}(0) = \bar{m}_0$. With such \bar{m} one solves the HJBI equation, first two lines of (4), and obtain a value function $v_{\bar{m}(\cdot)}$. The value function is then used to calculate the best-response u^* and worst adversarial disturbance w^* . These are plugged into the ODE of the mean distribution evolution, last two equations. We then solve the ODE and obtain a new function $\bar{M}_{\bar{m}(\cdot)}$. The process halts when $\bar{M}_{\bar{m}(\cdot)}(\cdot) = \bar{m}(\cdot)$.

A main issue in the above procedure is to prove in an abstract way that such a fixed point procedure leads to one and only one fixed point; a second issue is related to finding the value function $v_{\bar{m}(\cdot)}$ in order to have its derivative for the explicit calculation in the last line of (4). In this perspective a main question is the following one: does a solution of the first two lines necessarily have a spatial derivative? In the rest of the paper we address the above issues by using the weak notion of viscosity solutions.

III. MEAN FIELD EQUILIBRIUM

Consider a two player repeated game where player 1 plays $u(t) \in \Delta(U)$, player 2 plays $w \in \Delta(W)$, and $f : \Delta(U) \times$

$\Delta(W) \rightarrow [-1, 1]^n$ is the payoff. We assume that f is bilinear in u and w . We denote by G the one-shot vector payoff game $(\Delta(U), \Delta(W), f(t))$. With respect to the above game, in the spirit of attainability [14], we aim at analyzing convergence properties of the disturbed cumulative payoff $X(t)$ obtained by integrating (2) which we rewrite below

$$dX(t) = f(u(t), w(t))dt + \sigma dB(t), \quad X(0) = x.$$

Toward this aim, we will make use of the notion of *projected game* which we recall next. Let us consider $\lambda \in \mathbb{R}^n$ and denote by $\langle \lambda, G \rangle$ the one-shot zero-sum game whose set of players and their actions are as in game G , and the payoff that player 2 pays to player 1 is $\lambda^T f(u, w)$. Note that, as a zero-sum one-shot game, the game $\langle \gamma, G \rangle$ has a *value*, denoted $val[\lambda]$, obtained as

$$val[\lambda] := \inf_{u \in \Delta(U)} \sup_{w \in \Delta(W)} \{f(u, w)\lambda\}.$$

Assumption 1: The value of the projected game, $val[\lambda]$, is negative for every $\lambda \in \mathbb{R}^n$.

Now, for given x , take for λ the value $\lambda(v_x) = \frac{v_x(x, t)}{\|v_x(x, t)\|}$ which is the gradient direction on x . Then, we can introduce the value of the *projected anti-gradient game*

$$val[v_x(x, t)] := f(u^*, w^*)\lambda(v_x).$$

Due to the bilinear structure of f , we can deduce that the best-response strategy u^* and worst adversarial disturbance w^* are on a vertex of the associated simplices in \mathbb{R}^p and \mathbb{R}^q , respectively. This corresponds to saying that both strategies are *pure strategies*. We recall here that pure strategies are such that each player chooses a single action, in contrast with *mixed strategies* where players select probabilities on actions. A consequence of this is that the mean field equilibrium, if exists, is in pure strategies as well.

We can formalize the above concept by introducing

$$A_{ij} := f(\mathbf{1}_i, \mathbf{1}_j) \in [-1, 1]^n,$$

where $\mathbf{1}_i \in \mathbb{R}^p$ is a unit vector with all zero components except the i th one which is equal to 1, and similarly for $\mathbf{1}_j \in \mathbb{R}^q$. Then, we can rewrite the value of the anti-gradient projected game as

$$val[v_x] = \inf_{i \in I} \sup_{j \in J} A_{ij} \lambda(v_x),$$

where $I = \{1, \dots, p\}$ and $J = \{1, \dots, q\}$ are opportune sets of indices. Best responses and adversarial strategies are then

$$(i^*, j^*) = \arg \inf_{i \in I} \sup_{j \in J} A_{ij} \lambda(v_x).$$

With the above definition of $val[v_x]$ in mind, the Hamilton-Jacobi part of (4) can be rewritten as

$$\begin{aligned} v_t + \|v_x\| val[v_x] + \frac{1}{2} (\bar{m}(t) - x(t))^T Q (\bar{m}(t) - x(t)) \\ + \frac{\sigma^2}{2} v_{xx}^2 = 0 \text{ in } \mathbb{R}^n \times [0, T], \quad (5) \\ v(x, T) = \Psi(x) \quad \forall x \in \mathbb{R}^n. \quad (6) \end{aligned}$$

Our goal is to find a suitable final cost Ψ such that there exists a solution with $\bar{m} \equiv 0$. In order to do this, consider the solution Ψ of the stationary equation

$$\|\Psi_x\| = \frac{1}{-val[\Psi_x]} \left(\frac{1}{2} (\bar{m}(t) - x(t))^T Q (\bar{m}(t) - x(t)) + \frac{\sigma^2}{2} v_{xx}^2 \right).$$

We can show that, given function $\Psi_x(x)$ as above, any primitive $\hat{\Psi}(x)$ is a candidate value function for our problem.

Theorem 1: Let Assumption 1 hold. Then any primitive $\hat{\Psi}(x)$ is a candidate value function for our problem.

In order to arrive at the value function it remains to define Ω , the set of an associated exit-time optimal control problem yet to be introduced. Set Ω consists of all states x sufficiently close to zero, where a single player can drive the state to zero in one shot in the case of no Brownian motion. We can define such a set in two steps. First, let us identify the region of the state space where the adversary responds with a fixed pure strategy $l \in J$. This region is a cone with vertex at the origin, as the value function gradient v_x is radial and the adversary strategy depends on this only. More formally, for fixed $l \in J$ the cone we look for is

$$C_l = \{x \in \mathbb{R}^n \mid (i^*, j^*) = \arg \inf_{i \in I} \arg \sup_{j \in J} A_{ij} \lambda(v_x) = (i^*, l)\}.$$

For the same $l \in J$, we consider the set of states that can be driven to zero in one iteration. Such a set is a polyhedron and is defined as

$$\begin{aligned} P_l = -A_l^T \Delta(U) = \{x \in \mathbb{R}^n \mid -A_l^T u = xdt, \\ \text{for some } u \in \Delta(U)\}. \end{aligned}$$

Then, let us take the intersection of the two sets $\Omega_l = C_l \cap P_l$. The set of states for which a player can exit the game is then

$$\Omega = \cup_{l \in J} \Omega_l.$$

We are now in a position to formalize the exit time optimal control problem. Consider the function on \mathbb{R}^n

$$\Psi(x) = \begin{cases} \hat{\Psi}(x) & \text{if } x \notin \Omega, \\ \frac{1}{2} (\bar{m}(t) - x(t))^T Q (\bar{m}(t) - x(t)) \\ = \hat{\Psi}(x) & \text{if } x \in \Omega. \end{cases}$$

Also, consider the final-time problem in $\mathbb{R}^n \times [0, T]$ given by the running cost $g(\cdot, \cdot, 0)$ and by the continuous final cost Ψ . For every $\Omega \neq x \in \mathbb{R}^n$, we denote by $t^*(x)$ the reaching time of Ω under the nonanticipative minmax strategies (u^*, w^*) . It holds that the map $x \mapsto t^*(x)$ can be obtained through numerical computations and by solving the problem below in the variable $t^*(x)$ for a given x :

$$\begin{cases} \tilde{k} - x^T x = \int_0^{t^*(x)} (val[v_x] + \frac{\sigma^2}{2} v_{xx}^2) dt \\ dX(t) = f(u^*, w^*) dt + \sigma dB(t), \quad X(0) = x. \end{cases} \quad (7)$$

In the above system, we consider the balls centered at the origin with radius k , $B_k = \{x \in \mathbb{R}^n \mid x^T x \leq k\}$ and then compute the maximal inscribed ball $B_{\tilde{k}}$ in Ω

$$\tilde{k} := \arg \max\{k \mid B_k \subseteq \Omega\}.$$

Then, the value function of the optimal control problem introduced (the solution of the HJB equation) is given by

$$v(x, t) = \begin{cases} \Psi(x) & \text{if } T - t < t^*(x), \\ \Psi(x) + \frac{Q\sigma^2}{2}(T - t^*(x) - t) & \text{otherwise.} \end{cases}$$

In summary, the state of each single player reaches Ω in minimum time $t^*(x)$, and once there it oscillates around zero due to the Brownian motion. This justifies the cost $\frac{Q\sigma^2}{2}(T - t^*(x) - t)$ from $t^*(x)$ to T . Should $t^*(x)$ exceed the residual horizon length, the cost is then $\Psi(x)$.

In order to provide a macroscopic description of the system evolution let us start by discretizing the state space \mathcal{X} . Also let us denote by π the probability distribution vector defined on \mathcal{X} and let us refer to s as the total number of states, i.e., s is the dimension of \mathcal{X} , $\dim(\mathcal{X}) = s$. The FPK equation can then be replaced by the following Markov chain equation

$$\pi^+ = \pi P,$$

where P is a right (row) stochastic matrix in $\{0, a\}^{s \times s}$ and $a = 2^n$. If we zoom in on matrix P , then we notice that it is a block matrix of the form:

$$P = \begin{bmatrix} A & B \\ C & D \end{bmatrix}, \quad A = \left\{ \frac{1}{a} \right\}^{a \times a} \quad B = \{0\}^{a \times s} \\ C = \{0, a\}^{(s-a) \times a} \quad D = \{0, a\}^{(s-a) \times (s-a)}$$

Matrix A is responsible for stationarity, i.e., the stationary solution is a uniform distribution over the first a states, $[\pi_1 \dots \pi_a] = [1/a \dots 1/a]$. Such a solution satisfies the stationarity equation

$$[\pi_1 \dots \pi_a] = [\pi_1 \dots \pi_a]A.$$

In other words, once the system reaches $\mathcal{S} = \{\pm\sigma\}^n$, it never leaves the set. We also say that \mathcal{S} is *invariant*. For the same reason, matrix B is the zero matrix. A different way to see this is that A is row stochastic and therefore the rows of B must sum to zero. This observation, together with the fact that the elements of P are nonnegative suggests that all elements of B must be zero. Matrix C is a full column rank, as any point in \mathcal{S} is reachable from outside \mathcal{S} . Finally, matrix D has zero elements in the principal diagonal, as outside \mathcal{S} no states can keep their values fixed, or, saying the same, the transition probability from any given state to the same state is zero.

With the above in mind, we can infer that the probability distribution π converges to the stationary uniform distribution $[\pi_1 \dots \pi_a] = [1/a \dots 1/a]$. The same applies to m if we track back the discussion to its origins. The above distribution is zero-mean and variance equal to σ^2 , which is in accordance with the Brownian motion characteristics introduced earlier.

IV. CONCLUSIONS

This paper has shown how repeated games and population games can be intertwined to capture interactions among homogeneous social groups when herding behavior is rewarding for the groups. For the selected games, we have establish

a mean field equilibrium and studied state feedback best-response strategies as well as worst-case adversarial disturbances. Future directions of research involve the extension of the framework to other social behaviors (other types of cost functions), as well as social dynamics. The impact on social networks that existing results and techniques from repeated and population game literature can have is still a broad and open field.

REFERENCES

- [1] Y. Achdou, F. Camilli, I. Capuzzo Dolcetta, Mean field games: numerical methods for the planning problem, *SIAM J. of Control Opt.* vol. 50, pp. 77–109 (2012).
- [2] Y. Achdou and I. Capuzzo Dolcetta, Mean field games: numerical methods, *Siam J. Numer. Anal.*, vol. 48, pp. 1136–1162 (2010).
- [3] A. V. Banerjee, A simple model of herd behavior, *Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817 (1992).
- [4] M. Bardi, Explicit solutions of some linear-quadratic mean field games, *Network and Heterogeneous Media*, 7, pp. 243–261 (2012).
- [5] M. Bardi, I. Capuzzo Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, Birkhäuser, Boston (1997).
- [6] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*, SIAM Series in Classics in Applied Mathematics, Philadelphia (1999).
- [7] D. Bauso, H. Tembine, and T. Başar, Robust mean field games with application to production of an exhaustible resource, *Proceedings of 7th IFAC Symposium on Robust Control Design*, Aalborg, Denmark (2012).
- [8] D. Bauso and Q. Zhu and T. Başar, Mixed integer optimal compensation: decompositions and mean-field approximations, *Proceedings of 2012 American Control Conference*, Montreal, Montreal, Canada (2012).
- [9] O. Gueant, J. M. Lasry, and P. L. Lions, Mean field games and applications, *Paris-Princeton Lectures*, Springer, pp. 1–66 (2010).
- [10] M.Y. Huang, P.E. Caines, and R.P. Malhamé, Large population cost-coupled LQG problems with non-uniform agents: individual-mass behaviour and decentralized ϵ -Nash equilibria, *IEEE Trans. on Automatic Control*, vol. 52, no. 9, pp. 1560–1571 (2007).
- [11] B. Jovanovic, and R. W. Rosenthal, Anonymous sequential games. *Journal of Mathematical Economics*, vol. 17, pp. 77–87 (1988).
- [12] A. Lachapelle, J. Salomon, and G. Turinici, Computation of mean field equilibria in economics, *Math. Models Meth. Appl. Sci.*, vol. 20, pp. 1–22 (2010).
- [13] J.-M. Lasry and P.-L. Lions, Mean field games, *Jpn. J. Math.*, vol. 2, pp. 229–260 (2007).
- [14] E. Lehrer, E. Solan, and D. Bauso, Repeated games over networks with vector payoffs: the notion of attainability, *Proceedings of the Int. Conf. on NETWORK Games, CONTROL and OPTimization (NetGCooP 2011)*, 12–14 Oct. 2011, Paris.
- [15] R. Pesenti, and D. Bauso, Mean field linear quadratic games with set up costs, *Proceedings of the Int. Conf. on NETWORK Games, CONTROL and OPTimization (NetGCooP 2011)*, 12–14 Oct. 2011, Paris.
- [16] R. Selten, *Preispolitik der Mehrproduktenunternehmung in der statischen theorie*, Springer-Verlag, 1970.
- [17] H. Tembine, Mean field stochastic games, *Notes* (2011).
- [18] H. Tembine, Q. Zhu, and T. Başar, Risk-sensitive mean-field stochastic differential games, *Proceedings of 2011 IFAC World Congress*, Milan, Italy, August 29 - September 2, 2011.
- [19] H. Tembine, J. Y. Le Boudec, R. ElAzouzi, E. Altman, Mean field asymptotic of Markov decision evolutionary games, *International IEEE Conference on Game Theory for Networks*, Gamenets 2009.
- [20] Q. Zhu, H. Tembine, and T. Başar, Hybrid risk-sensitive mean-field stochastic differential games with application to molecular biology, *Proceedings of Conference on Decision and Control*, Orlando, FL, Dec. 14 - Dec. 16, 2011.
- [21] Q. Zhu and T. Başar, A multi-resolution large population game framework for smart grid demand response management, *Proceedings of Intl. Conference on Network Games, Control and Optimization (NETGCOOP 2011)*, Paris, France. Oct. 12 - 14, 2011.