Synchronization Games*

Felix Höfer and H. Mete Soner †

February 23, 2024

Abstract

We propose a new mean-field game model with two states to study synchronization phenomena, and we provide a comprehensive characterization of stationary and dynamic equilibria along with their stability properties. The game undergoes a phase transition with increasing interaction strength. In the subcritical regime, the uniform distribution, representing incoherence, is the unique and stable stationary equilibrium. Above the critical interaction threshold, the uniform equilibrium becomes unstable and there is a multiplicity of stationary equilibria that are self-organizing. Under a discounted cost, dynamic equilibria spiral around the uniform distribution before converging to the self-organizing equilibria. With an ergodic cost, however, unexpected periodic equilibria around the uniform distribution emerge.

Keywords: Kuramoto synchronization, mean-field games, consensus problems, Markov processes, Nash equilibrium, dynamic programming.

Mathematics Subject Classification: 34C25, 34H05, 37G35, 49L20, 91A16, 92B25

1 Introduction

Building on Winfree's work in the 1960s, the Kuramoto model [27] has become the corner stone of mathematical models of collective synchronization and has received attention in all natural sciences, engineering, and mathematics. The model consists of coupled oscillators that exhibit spontaneous synchronization once the coupling strength exceeds a critical threshold. While the classical model postulates the dynamics of each oscillator in the form of a system of nonlinear ordinary differential equations, we use the mean-field game (MFG) formalism. Indeed, instead of positing the dynamics of the particles, we let the individual particles determine their behavior endogenously by minimizing a cost functional and settling in a Nash equilibrium. More generally, MFGs were independently introduced by Lasry & Lions [28]–[30] and Huang, Caines, & Malhamé [20]–[23] to approximate large population games in which the interaction appears through the empirical distribution of all agents.

The MFG approach to synchronization was first proposed by Yin, Mehta, Meyn, & Shanbhag in [33], [34], and later used by Carmona & Graves [7] to study jet-lag recovery by modeling the alignment with the circadian rhythm. These studies and the recent work of Carmona, Cormier, & Soner [5] establish phase transitions in the Kuramoto MFGs analogous to the one exhibited by the original Kuramoto dynamical system, providing strong evidence for a connection between these two seemingly very different models.

In the mean-field game, each oscillator is treated as a rational agent that minimizes the distance to other oscillators while incurring a quadratic cost. The coupling strength, or the misalignment cost, is the central parameter in these models and the phase transition can be summarized as follows:

- *Subcritical regime*: Below the critical coupling strength, oscillators behave incoherently and the uniform distribution is the unique stationary Nash equilibrium. Additionally, [5] obtains a local stability result by proving that from any initial condition that is sufficiently close to the uniform distribution there exist time-inhomogeneous Nash equilibria converging to the uniform measure as time passes.
- *Supercritical regime*: Above the critical interaction parameter, multiple non-uniform stationary Nash equilibria emerge, leading to partial self-organization. As the interaction strength goes to infinity, these stationary Nash equilibria become more coherent and converge to a Dirac measure representing full synchronization.

*This work was partially supported by the National Science Foundation grant DMS 2106462.

†Both authors are with Princeton University, Department of Operations Research and Financial Engineering (e-mails: fhoefer@princeton.edu, soner@princeton.edu).

Recently, Cesaroni & Cirant [12] obtained further results with an ergodic cost functional in the *large* parameter regime. They prove that for large misalignment parameters, there is only one non-uniform, self-organizing, stationary Nash equilibrium modulo translations. Furthermore, [12] show that there are Nash equilibria of the finite horizon problem that converge to the self-organizing equilibrium.

In the Kuramoto model, the one-dimensional torus given by $\mathbb{T} = \mathbb{R}/(2\pi\mathbb{Z})$ represents the set of possible phases of the oscillators. We simplify this by considering an equidistant discretization $\mathcal{X}_N = \{x_1, \dots, x_N\} \subset \mathbb{T}$ of the phase space, as done more generally by Bertucci & Cecchin [4]. We assume that the oscillators are subject to non-zero *thermal noise* of strength $\sigma^2 > 0$ and without any control, they would move to their left or right on $\mathcal{X}_N$ with a rate of $N^2\sigma^2/2$. In the game setting, oscillators are allowed to influence this transition by choosing a feedback control $\alpha := (\alpha_\ell, \alpha_r) : [0,\infty) \times \mathcal{X}_N \to [0,\infty)^2$. Then, the position of a generic oscillator is described by a continuous-time Markov chain X_t^α with controlled transition rates $\lambda_{i,j}^\alpha = \lambda(x_i, x_j, \alpha)$ from state x_i to x_j given by,

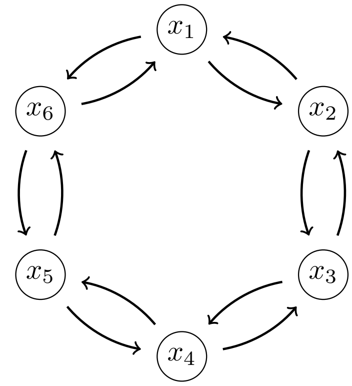

Figure 1: Example of a six-state model.

$$\lambda_{i,j}^\alpha = \begin{cases} N\alpha_\ell(t, x_i) + N^2\sigma^2/2 & \text{if } j = i-1 \ (\mathrm{mod}\ N), \\ N\alpha_r(t, x_i) + N^2\sigma^2/2 & \text{if } j = i+1 \ (\mathrm{mod}\ N), \\ 0 & \text{else.} \end{cases}$$

Figure 1 is a visualization of a six-state model.

We now describe the *discretized Kuramoto MFG* which is a game between infinitely many such oscillators while the number of possible phases N is kept finite. Let the probability flow $(\mu_t)_{t\geq 0}$ on $\mathcal{X}_N$ represent the population distribution of oscillators' phases. Then, μ_t solves a forward Kolmogorov equation which the representative oscillator or agent cannot influence. Instead, they try to align their own phase with μ_t by minimizing the corresponding discounted infinite cost,

$$J_\beta(\alpha) := \mathbb{E} \int_0^\infty e^{-\beta t} \left(\frac{1}{2} |\alpha(t, X_t^\alpha)|^2 + \kappa\, \ell(X_t^\alpha, \mu_t) \right) \mathrm{d}t,$$

over feedback controls α, where $X_t^\alpha \in \mathcal{X}_N$ is the random position of the representative oscillator with controlled rates $(\lambda_{i,j}^\alpha)$ and $X_0^\alpha \sim \mu_0$, the discount factor is $\beta > 0$, the running cost ℓ is defined by

$$\ell(x, \mu) := 2 \sum_{y \in \mathcal{X}_N} \sin^2\left(\frac{x-y}{2}\right) \mu(\{y\}),$$

and $\kappa > 0$ is the strength of the interactions between oscillators. We emphasize again that μ_t is fixed for the representative agent, and any optimal control α^* and the distribution $\mathcal{L}(X_t^{\alpha^*})$ is a function of μ_t. Hence, this construction defines a map $(\mu_t)_{t\geq 0} \mapsto (\mathcal{L}(X_t^{\alpha^*}))_{t\geq 0}$, and the fixed points of this map are the *mean-field game (Nash) equilibria*. If a Nash equilibrium $t \mapsto \mu_t^*$ is constant, we call it a *stationary MFG (Nash) equilibrium*. We say that a stationary Nash equilibrium μ^* is *stable* if all time-inhomogeneous Nash equilibria $(\mu_t^*)_{t\geq 0}$ converge towards μ^* as $t \to \infty$.

1.1 Results

We study a two-state model, providing a complete characterization of the set of stationary and time-inhomogeneous equilibria and their stability properties for all values of the parameters $(\beta, \sigma^2, \kappa)$. In addition to the discounted cost functional, we also investigate the models with ergodic cost which formally correspond to the limit $\beta \downarrow 0$, and further identify the equilibria that are obtained as first-order conditions of the associated mean-field control problem, see Sections 5.3 and 6.1. A short discussion of the finite horizon problem and its convergence to the ergodic one is given in subsection 6.2.

In what follows, the state space $\mathcal{X}_2$ is mapped to the discrete set $\mathcal{X} := \{0, 1\}$, and a probability measure μ on $\mathcal{X}$ is identified with its value $p = \mu(\{1\})$. For a measurable real-valued function $\alpha : [0,\infty) \to \mathbb{R}$ we construct a unique feedback control by letting the control from state 0 to 1 be $\alpha(t, 0) := \alpha^+(t)$ and the control from 1 to 0 be $\alpha(t, 1) := \alpha^-(t)$. These special feedback controls are the only ones that satisfy $\alpha(t, 0)\alpha(t, 1) = 0$ for all $t \geq 0$. In the two state model, it is clear that optimal controls also have this property. Therefore, there is no loss of generality to consider the set $\mathcal{A}$ of measurable functions $\alpha : [0,\infty) \to \mathbb{R}$ to be the *feedback controls*, so that $\sigma^2 + \alpha^+(t)$ is the rate from 0 to 1 and $\sigma^2 + \alpha^-(t)$ is the rate from 1 to 0.

Suppose that a flow of probabilities $p(\cdot)$ is a Nash equilibrium. Then, there is an optimal feedback control $a \in \mathcal{A}$ such that $p(t) = \mathbb{P}(X_t^a = 1)$. Consequently, these probabilities solve the forward Kolmogorov equation,

$$\dot{p}(t) = (\sigma^2 + a^+(t))(1 - p(t)) - (\sigma^2 + a^-(t))p(t). \tag{1}$$

Moreover, by standard techniques from optimal control, we show in subsection 3.3 that for a given flow of probabilities $p(\cdot)$ there is a unique optimal feedback control $a(\cdot)$ solving

$$\dot{a}(t) = (\beta + 2\sigma^2)a(t) + \frac{1}{2}\mathrm{sign}(a(t))a(t)^2 - \kappa(2p(t) - 1), \tag{2}$$

where $\beta = 0$ corresponds to the ergodic cost. Combining these, we characterize all Nash equilibria as bounded solutions of (1,2). This one-to-one connection between the Nash equilibria and (1,2) is established in Propositions 3.1 and 3.2. We continue by summarizing its consequences, while precise statements and their proofs are provided in the following sections.

1.1.1 Discounted cost

For a positive discount factor $\beta > 0$ the critical coupling or misalignment strength is

$$\kappa_c = 2\beta\sigma^2 + 4\sigma^4.$$

Subcritical regime $\kappa < \kappa_c$. The uniform distribution is the unique stationary Nash equilibrium (SNE) and it is stable. Further, for any initial distribution there exists a unique time-inhomogeneous Nash equilibrium (NE), and it converges to the uniform distribution.

Supercritical regime $\kappa > \kappa_c$. There exist three SNE: the uniform distribution and two symmetric self-organizing SNE given by $\underline{p} < 1/2 < \overline{p}$ where $\underline{p} = 1 - \overline{p}$. The local behavior around the uniform distribution depends on two sub-regimes:

(A) $\kappa_c < \kappa < \kappa_c + \beta^2/4$. For any non-uniform initial distribution there exists a unique NE and it converges to one of the self-organizing SNE.

(B) $\kappa > \kappa_c + \beta^2/4$. For initial conditions close to the uniform distribution there exist many NE that spiral around the uniform distribution before converging to one of the self-organizing SNE.

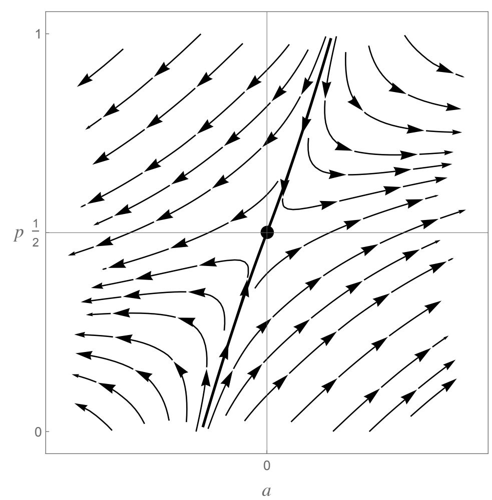

Figure 2: Phase diagram of the discounted subcritical system (1,2). The dot shows the location of the stationary Nash equilibrium while the thick lines illustrate time-inhomogeneous Nash equilibria.

We visualize these results by the phase diagrams of the ordinary differential equation (1,2). The uniform SNE with zero control corresponds to the fixed point $(a, p) = (0, 1/2)$ of this system of equations. The local behavior of the solutions around this point changes with increasing interaction and can be summarized by the eigenvalues λ_1, λ_2 of the linearized system around $(0, 1/2)$. In the subcritical regime, the point $(0, 1/2)$ is a saddle, i.e. $\lambda_1 < 0 < \lambda_2$, and its stable manifold crosses the boundary $\{p = 0, 1\}$ as seen in Figure 2.

When $\kappa > \kappa_c$, the two self-organizing SNE are saddles and their stable manifolds join at the origin, creating a curve $\mathcal{C}$ that connects all SNE and hits the boundaries $\{p = 0, 1\}$. In the weakly supercritical regime (A), the curve $\mathcal{C}$ is monotone and the uniform SNE becomes repellent with both eigenvalues being positive, as shown in Figure 3a. In the supercritical regime (B) however, λ_1, λ_2 become complex with positive real parts, so that $(0, 1/2)$ is a spiral source, see Figure 3b. In all phase diagram, the thick dots show the location of stationary equilibria.

All Kuramoto MFGs are potential and there is an associated mean-field control (MFC) problem introduced in subsections 5.3, 6.1 below. The minimizers of these problems are Nash equilibria for the MFG. Surprisingly, with a discounted cost, the uniform distribution is selected as the minimizer even in parts of the supercritical regime.

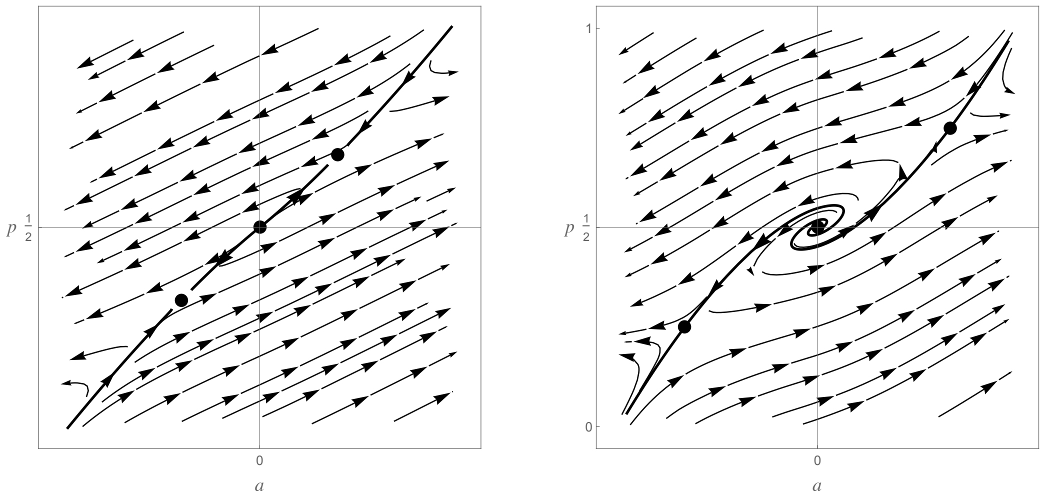

(a) Supercritical case (A): $\kappa_c < \kappa < \kappa_c + \beta^2/4$. The thick line corresponds to the monotone curve $\mathcal{C}$, and the dots show the stationary equilibria.

(b) Supercritical case (B): $\kappa > \kappa_c + \beta^2/4$. The thick line again corresponds to the curve $\mathcal{C}$, which spirals around $(0, 1/2)$.

Figure 3: Phase diagrams of the supercritical discounted system (1,2)

1.1.2 Ergodic cost

Here the critical interaction parameter is

$$\kappa_c = 4\sigma^4.$$

In the subcritical regime, the uniform distribution $p^* = 1/2$ is again the unique SNE, and there are no other ergodic NE.

In the supercritical case $\kappa > \kappa_c$, in addition to the uniform distribution, there are two other symmetric SNE. However, in contrast to the discounted model, there are infinitely many periodic NE rotating around the uniform distribution as well. Figure 4 shows the phase diagram of the system (1,2) close to the uniform distribution, and in dynamical systems terminology, the origin is a center of the equations (1,2).

In the subcritical regime, the uniform distribution is the minimizer of the associated ergodic mean-field control (MFC) problem. It fails to be the minimizer in the supercritical case.

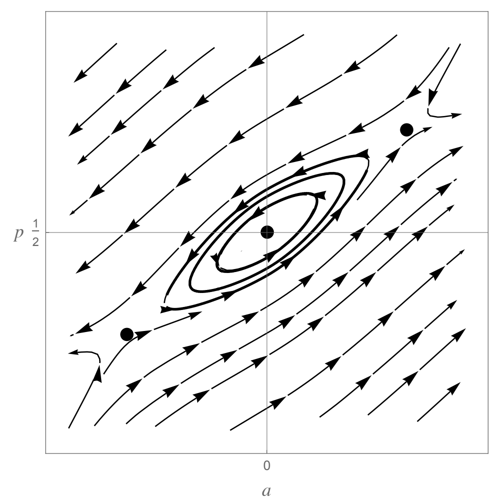

Figure 4: Phase diagram of the supercritical ergodic system (1,2).

1.2 Related studies

This paper studies the emergence and properties of self-organizing equilibria as the outcome of a game among a continuum of rational agents that favor alignment with the majority. As such it is situated at the intersection of several fields. In addition to several already mentioned studies on Kuramoto MFG and synchronization, our model is related to the studies of opinion dynamics as studied in [1], [3], [9] and the references therein. The static mean-field game *"Where do I put my towel on the beach?"* that Lions discusses in [31] is also similar to our model. In such a setting, there is non-uniqueness once agents favor crowds as opposed to avoiding them. Also, Dai Pra, Sartori, & Tolotti [14] investigate a dynamic model that is similar to ours in the special case of $\sigma = 0$ to study the emergence of collective behavior. The absence of thermal noise, however, precludes the emergence of the phase transition and random structures.

Furthermore, our work contributes to a growing literature of finite-state mean-field games [6][Chapter 7.2]. These include [24]–[26] who study the spread of corruption and a botnet defense model. The first two papers analyze a game with three and four states in a bang-bang type model that exhibits a multiplicity of stationary equilibria. Socio-economic applications of finite-state MFGs, including a description of the potential game structure, can be found in [15]. Classical results establishing existence, uniqueness under the Lasry-Lions monotonicity condition, and the convergence of the finite player games in finite-state MFGs can for example be found in [8], [10], [17]. Finally, Cohen & Zell [13] treat the infinite horizon case, and [11] studies the convergence problem of a two-state MFG under an anti-monotonous cost.

2 Two state mean-field game

Recall that $\mathcal{X} = \{0, 1\}$ and we identify any $p \in [0, 1]$ with a measure on $\mathcal{X}$ whose value on $\{1\}$ is equal to p. For $x \in \mathcal{X}$, $x + 1$ denotes summation mod 2. The set of feedback controls $\mathcal{A}$ consists of measurable functions $\alpha : [0, \infty) \to \mathbb{R}$ that are locally integrable. Following the interpretation discussed in the Introduction, for any $\alpha \in \mathcal{A}$ and initial distribution there exists a Markov chain X^α_t with the transition matrix

$$Q^\alpha_t = \begin{pmatrix} -\sigma^2 - \alpha^+(t) & \sigma^2 + \alpha^+(t) \\ \sigma^2 + \alpha^-(t) & -\sigma^2 - \alpha^-(t) \end{pmatrix},$$

where $\sigma^2 > 0$ is the strength of the thermal noise. For $x \in \mathcal{X}$, $p \in [0, 1]$, we introduce the *running cost* ℓ by

$$\ell(x, p) := \begin{cases} p & \text{if } x = 0, \\ 1 - p & \text{if } x = 1. \end{cases}$$

Further, we denote the misalignment or the coupling strength by $\kappa > 0$, and call it the *coupling constant*.

Definition 2.1 *Fix a discount factor $\beta > 0$ and a coupling constant κ. We say that a flow of probabilities $t \mapsto p(t) \in [0, 1]$ is a* discounted mean-field game Nash equilibrium *starting from $p(0)$, if there exists $\alpha^* \in \mathcal{A}$ such that*

1. *α^* minimizes the discounted cost functional*
$$J_\beta(\alpha) := \mathbb{E} \int_0^\infty e^{-\beta t} \left(\frac{1}{2} \alpha(t)^2 + \kappa \ell(X^\alpha_t, p(t)) \right) \mathrm{d}t,$$
where X^α_t is the inhomogeneous Markov chain with $\mathbb{P}[X^\alpha_0 = 1] = p(0)$ and transition matrix $(Q^\alpha_t)_{t \geq 0}$.
2. *For all times $t \geq 0$, we have $p(t) = \mathbb{P}[X^{\alpha^*}_t = 1]$.*

We next consider the ergodic cost.

Definition 2.2 *We say that a* periodic *or* stationary *flow of probabilities $p(\cdot) \in [0, 1]$ is an* ergodic mean-field game Nash equilibrium *if there exists $\alpha^* \in \mathcal{A}$ such that*

1. *α^* minimizes the ergodic cost functional*
$$J_e(\alpha) := \overline{\lim_{T \uparrow \infty}} \frac{1}{T} \mathbb{E} \int_0^T \left(\frac{1}{2} \alpha(t)^2 + \kappa \ell(X^\alpha_t, p(t)) \right) \mathrm{d}t.$$
2. *For all times $t \geq 0$, we have $p(t) = \mathbb{P}[X^{\alpha^*}_t = 1]$.*

As $t \mapsto \mathbb{P}(X^\alpha_t = 1)$ is continuous for any $\alpha \in \mathcal{A}$, all NE $p(\cdot)$ are continuous.

Since $J_e(\alpha)$ is invariant under changes of the control and the probabilities on bounded time intervals, one has to restrict either the controls or the probabilities to obtain a meaningful definition. In the above, we use the periodicity condition on the probabilities since it does not cause a loss of generality in our model. Indeed, as we will show in the next subsection, all Nash equilibria are characterized by a planar dynamical system and the limit behavior of such systems is always periodic.

3 Characterization of Nash Equilibria

Nash equilibria are in a natural one-to-one correspondence with a system of forward-backward differential equations which we refer to as the *MFG system*. To state this characterization, for $x \in \mathcal{X}$, $v \in \mathbb{R}$, and $p \in [0, 1]$, we define a *Hamiltonian* by

$$H(x, v, p) := \inf_{z \geq 0} \left\{ (\sigma^2 + z)v + \frac{1}{2} z^2 \right\} + \kappa\, \ell(x, p) = \sigma^2 v - \frac{1}{2}(v^-)^2 + \kappa \ell(x, p). \tag{3}$$

3.1 Ergodic MFG system

The ergodic dynamic programming equation is

$$-v_t(t,x) + \overline{\lambda} = H(x, v(t,x+1) - v(t,x), p(t)), \tag{4}$$

for every $x \in \mathcal{X}, t \geq 0$, and the unique optimal control is given by

$$a(t) = v(t,0) - v(t,1). \tag{5}$$

Classically, given a flow $p(\cdot)$, such a solution pair $(\overline{\lambda}, v)$ is constructed by letting β to zero in the corresponding discounted control problem. Indeed, if v_β is the value function of the discounted problem, then $\overline{\lambda}$ is the limit of βv_β, and $v(t,x)$ is the limit of $v_\beta(t,x) - v_\beta(0,0)$, see [2].

When the flow $p(\cdot)$ is a Nash equilibrium, then $p(\cdot)$ solves (1) with the feedabck control defined above, and the ergodic mean-field game system (MFG_0) consists of the equations (4), (1) coupled by (5). Precisely, we call a triplet $(\overline{\lambda}, v, p)$ a *classical solution* of (MFG_0) if $\overline{\lambda} \in \mathbb{R}$ is a real number, and $v : [0,\infty) \times \mathcal{X} \to \mathbb{R}$, $p : [0,\infty) \to [0,1]$ are continuously differentiable in the time variable, and they satisfy (4) and (1) with a given by (5).

We proceed to establish the correspondence of NE with solutions of (MFG_0), and also show that $\overline{\lambda}$ has an interpretation as the optimal value of the ergodic cost,

$$\overline{\lambda} = \inf_{\alpha \in \mathcal{A}} \ \overline{\lim_{T\uparrow\infty}} \frac{1}{T} \mathbb{E} \int_0^T \left(\frac{1}{2} \alpha(t)^2 + \kappa \ell(X_t^\alpha, p(t)) \right) \mathrm{d}t. \tag{6}$$

Let $\mathcal{A}_c$ be the set of feedback controls in $\mathcal{A}$ that are continuous in the time variable.

Proposition 3.1 *For a periodic flow of probabilities $p(\cdot)$, the following are equivalent:*

(i) $p(\cdot)$ is an ergodic Nash equilibrium.

(ii) There are a constant $\overline{\lambda}$ and $v : [0,\infty) \times \mathcal{X} \to \mathbb{R}$ such that $(\overline{\lambda}, v, p)$ is a bounded classical solution of (MFG_0).

(iii) $p(\cdot)$ is an ergodic Nash equilibrium with an optimal feedback control $a \in \mathcal{A}_c$.

(iv) There is a feedback control $a \in \mathcal{A}_c$ such that (a, p) is a periodic solution of (1,2) with $\beta = 0$.

In all cases, $\overline{\lambda}$ is given by (6).

Proof. (*ii*) $\Rightarrow$ (*iii*). Since $v(\cdot, x)$ is continuously differentiable, Dynkin's formula implies that for any feedback control $\alpha \in \mathcal{A}$,

$$\mathbb{E}[v(t, X_T^\alpha)] - \mathbb{E}[v(0, X_0^\alpha)] = \mathbb{E} \int_0^T (\partial_t + A^\alpha) \, v(t, X_t^\alpha) \, \mathrm{d}t \geq T\overline{\lambda} - \mathbb{E} \int_0^T \left(\frac{1}{2} \alpha(t)^2 + \kappa \ell(X_t^\alpha, p(t)) \right) \mathrm{d}t,$$

where $X_0^\alpha \sim p(0)$ and A^α is the infinitesimal generator of the controlled Markov chain: for $x \in \mathcal{X}$, $t \geq 0$,

$$A^\alpha \Phi(t,x) = (\sigma^2 + \alpha(t,x))(\Phi(t,x+1) - \Phi(t,x)),$$

where with an abuse of notation, we set $\alpha(t,0) := \alpha^+(t)$, $\alpha(t,1) := \alpha^-(t)$. Since v is bounded, dividing by T and sending $T \uparrow \infty$ shows that $\overline{\lambda}$ is a lower bound for the ergodic cost. Same argument with $a \in \mathcal{A}_c$ given by (5) shows (6). Hence, a is optimal. Since $\mathbb{P}[X_t^a = 1]$ satisfies the same forward Kolmogorov equation (1) as $p(\cdot)$, by uniqueness $p(t) = \mathbb{P}[X_t^a = 1]$ for all $t \geq 0$. Together with the periodicity condition, this proves that $p(\cdot)$ is an ergodic NE.

(*iii*) $\Rightarrow$ (*i*). This is trivial.

(*i*) $\Rightarrow$ (*ii*). Suppose $p(\cdot) = \mathbb{P}[X_t^{\alpha^*} = 1]$ is an ergodic NE corresponding to the optimal control $\alpha^* \in \mathcal{A}$. Let $(\overline{\lambda}, v(t,x))$ be a bounded solution of the dynamic programming equation (4) with $p(t)$, and define a by (5). We claim that for Lebesgue a.e. $t \geq 0$, $\alpha^*(t) = a(t)$. As $p(t) = \mathbb{P}[X_t^{\alpha^*} = 1]$ solves the Kolmogorov equation (1) with $\alpha^*(t)$ replacing $a(t)$, this claim would imply that the triplet $(\overline{\lambda}, v, p)$ is a bounded classical solution of (MFG_0). So it suffices to prove this claim.

Indeed, we first observe that since $p(\cdot)$ is periodic, so is the map $t \mapsto v(t,x)$. Let $\tau > 0$ be the common period, and let X_t^* be the Markov chain starting in $X_0^* \sim p(0)$ and controlled by α^*. Then, as $\dot{p}(\cdot)$ is τ-periodic, so is $\alpha^*(\cdot)$, and we obtain

$$\overline{\lim_{T\uparrow\infty}} \frac{1}{T} \ \mathbb{E} \int_0^T \left(\frac{1}{2} (\alpha^*(t))^2 + \kappa \ell(X_t^*, p(t)) \right) \mathrm{d}t = \frac{1}{\tau} \mathbb{E} \int_0^\tau \left(\frac{1}{2} (\alpha^*(t))^2 + \kappa \ell(X_t^*, p(t)) \right) \mathrm{d}t.$$

Towards a contraposition, assume there is $x' \in \mathcal{X}$, $\delta > 0$, and a subset $I \subset [0, \tau)$ of positive Lebesgue measure satisfying

$$(\sigma^2 + \alpha^*(t,x'))(v(t,1+x') - v(t,x')) + \frac{1}{2}(\alpha^*(t,x'))^2 + \kappa\ell(x',p(t)) - \delta \\ \geq H(x', v(t,1+x') - v(t,x'), p(t)), \qquad \forall t \in I.$$

By periodicity, equation (4), and the Dynkin's formula,

$$0 = \mathbb{E}[v(\tau, X^*_\tau)] - \mathbb{E}[v(0, X^*_0)] = \mathbb{E}\int_0^\tau (\partial_t + A^{\alpha^*})\, v(t, X^*_t)\,\mathrm{d}t \\ \geq \tau\bar{\lambda} - \mathbb{E}\int_0^\tau \left(\frac{1}{2}(\alpha^*(t))^2 + \kappa\ell(X^*_t, p(t))\right)\mathrm{d}t + \delta\int_I \mathbb{P}(X^*_t = x')\,\mathrm{d}t.$$

Hence, $\tau\bar{\lambda} + \delta\int_0^\tau \mathbb{P}(X^*_t = x')\,\mathrm{d}t \leq \tau\bar{\lambda}$. This is impossible, as $\sigma > 0$ precludes the Markov chain X^* to be identically equal to one state. This proves the claim.

(*iv*) $\Rightarrow$ (*ii*). Let (a, p) be a periodic solution of (1,2) with period $\tau > 0$. For any $\lambda \in \mathbb{R}$, and $x \in \mathcal{X}$, $t \in [0, \tau]$ set

$$V(t,0) := a(0) + \int_0^t [\lambda - H(0, -a(s), p(s))]\,\mathrm{d}s, \quad V(t,1) := \int_0^t [\lambda - H(1, a(s), p(s))]\,\mathrm{d}s,$$

and $A(t) := V(t,0) - V(t,1)$. We now use the fact that a satisfies (2) and the explicit form of H, to compute

$$\dot{A}(t) = H(1, a(t), p(t)) - H(0, -a(t), p(t)) = 2\sigma^2 a(t) - \frac{1}{2}(a(t)^-)^2 + \frac{1}{2}((-a(t))^-)^2 - \kappa(2p(t) - 1) \\ = 2\sigma^2 a(t) + \frac{1}{2}\mathrm{sign}(a(t))a(t)^2 - \kappa(2p(t) - 1) = \dot{a}(t).$$

As $A(0) = a(0)$, $A \equiv a$. Therefore, V solves the dynamic programming equation (4) with $p(\cdot)$. Since by hypothesis $p(\cdot)$ solves (1) with a, the triplet (λ, V, p) is a classical solution of (MFG$_0$) for any λ. To obtain a bounded solution, we choose

$$\lambda = \frac{1}{\tau}\int_0^\tau H(1, a(s), p(s))\mathrm{d}s.$$

Since a and p are τ-periodic, the above choice of λ ensures that $V(\cdot, 1)$ is also τ-periodic. Since $V(\cdot, 0) = V(\cdot, 1) + A(\cdot)$, it is also periodic. Therefore, V is periodic, and hence bounded.

(*ii*) $\Rightarrow$ (*iv*). Let $(\bar{\lambda}, v, p)$ be a bounded, classical solution of (MFG$_0$), and let a be as in (5). Then, a direct calculation shows that the pair (a, p) solves (1,2). $\square$

3.2 Discounted MFG system

Given a discount factor $\beta > 0$, a flow of probabilities $p(\cdot)$, $t \geq 0$, and $x \in \mathcal{X}$, the optimal control problem of the representative oscillator is given by

$$v(t,x) := \inf_{\alpha \in \mathcal{A}} \mathbb{E}\int_t^\infty e^{\beta(t-u)}\left(\frac{1}{2}\alpha(u)^2 + \kappa\ell(X^\alpha_u, p(u))\right)\mathrm{d}u,$$

where X^α is as before and $X^\alpha_t = x$. Let H be as in (3). Then, the dynamic programming equation for this problem is

$$-v_t(t,x) + \beta v(t,x) = H(x, v(t, x+1) - v(t,x), p(t)), \tag{7}$$

for all $x \in \mathcal{X}$, $t \geq 0$. Moreover, the optimal control a is given by (5). Then, in view of Definition 2.1, $p(\cdot)$ is a discounted NE if it solves (1) with this control a.

We say that a pair (v, p) is a *classical solution* of (MFG$_\beta$) with initial condition $p(0) \in [0, 1]$ if v is a classical solution of the dynamic programming equation (7) and p is classical solution of (1) with a given by (5). As for the ergodic cost we have the following characterization.

Proposition 3.2 *For $\beta > 0$ and a flow of probabilities $p(\cdot)$, the following are equivalent:*

(i) $p(\cdot)$ is a discounted Nash equilibrium.

(ii) There is $v : [0, \infty) \times \mathcal{X} \to \mathbb{R}$ such that (v, p) is a bounded classical solution of (MFG$_\beta$).

(iii) $p(\cdot)$ is a discounted Nash equilibrium with an optimal feedback control $a \in \mathcal{A}_c$.

(iv) There is a feedback control $a \in \mathcal{A}_c$ such that the pair (a, p) is a bounded solution of (1,2).

Proof. For all implications except $(iv) \Rightarrow (ii)$, we follow the proof of Proposition 3.1 *mutatis mutandis.*

$(iv) \Rightarrow (ii)$. Let (a,p) be a bounded solution of (1,2). We set

$$v(t,0) := \int_t^\infty e^{\beta(t-u)} H(0,-a(s),p(s))\ \mathrm{d}s, \qquad v(t,1) := \int_t^\infty e^{\beta(t-u)} H(1,a(s),p(s))\ \mathrm{d}s,$$

and $A(t) := v(t,0) - v(t,1)$. We directly show that

$$\dot{A}(t) = \dot{a}(t) + \beta(A(t) - a(t)), \quad \Rightarrow \quad \frac{\mathrm{d}}{\mathrm{d}t}(e^{-\beta t}(A(t) - a(t))) = 0, \quad \Rightarrow \quad e^{-\beta t}(A(t) - a(t))) = A(0) - a(0).$$

Since A and a are bounded, we conclude that $a(t) = A(t) = v(t,0) - v(t,1)$. Therefore,

$$-v_t(t,0) + \beta v(t,0) = H(0,-a(t),p(t)) = H(0, v(t,1) - v(t,0), p(t),$$
$$-v_t(t,1) + \beta v(t,1) = H(1,a(t),p(t)) = H(1, v(t,0) - v(t,1), p(t).$$

Hence, the pair (v,p) is a classical solution of (MFG$_\beta$).

□

3.3 Change of variables

We have shown that NE are characterized by bounded classical solutions to the coupled system of ordinary differential equations (1) and (2). For further analysis, it is convenient to introduce the following change of variables to achieve symmetry,

$$q(t) = 2p(t) - 1.$$

Then, (1,2) is equivalent to

$$\dot{a}(t) = (\beta + 2\sigma^2)a(t) + \frac{1}{2}\mathrm{sign}(a(t))a(t)^2 - \kappa q(t), \tag{2}$$
$$\dot{q}(t) = a(t) - (2\sigma^2 + |a(t)|)q(t). \tag{8}$$

We study this differential equation on the strip $(a,q) \in D := \mathbb{R} \times [-1,+1]$. Notice that D is invariant under the above equations.

4 Stationary Equilibria

A *stationary mean-field game Nash equilibrium* (SNE) is a constant mean-field game equilibrium. We should emphasize that in this case the initial condition $p(0)$ (or $q(0)$) is not given anymore, but becomes part of the solution. In view of the results of the Section 3 and the above change of variables, SNE are given by the second component of the fixed points of the planar system of ordinary differential equations (2, 8), both in the ergodic and discounted cases. To compute them, let S_a and S_q denote the nullclines of this system:

$$S_a = \{(a,q) \,|\, (\beta + 2\sigma^2)a + \mathrm{sign}(a)a^2/2 - \kappa q = 0\}, \qquad S_q := \{(a,q) \,|\, a - (2\sigma^2 + |a|)q = 0\},$$

and set

$$\kappa_c := 2\beta\sigma^2 + 4\sigma^4.$$

Clearly, $S_a \cap S_q$ are the fixed points of (2, 8), and the origin is always in this set. A direct calculation shows that $S_a \cap S_q$ is a singleton for $\kappa \le \kappa_c$. However, if $\kappa > \kappa_c$, there are three stationary points $S_a \cap S_q = \{(-\bar{a}, -\bar{q}), (0,0), (\bar{a}, \bar{q})\}$, where

$$\bar{q} = \frac{\bar{a}}{\bar{a} + 2\sigma^2}, \quad \bar{a} = -(\beta + 3\sigma^2) + \sqrt{(\beta + 3\sigma^2)^2 + 2(\kappa - \kappa_c)} > 0. \tag{9}$$

Hence, we have the following immediate corollary.

Lemma 4.1 *Both in the ergodic and discounted models, the uniform distribution $q = 0$ is a SNE for any $\kappa > 0$. Moreover,*

(i) Subcritical case: For $\kappa < \kappa_c$, the uniform distribution is the only SNE.

(ii) Supercritical case: For $\kappa > \kappa_c$, there are three SNE given by $q = -\bar{q}$, 0, $\bar{q}$.

(iii) Full synchronization: As $\kappa \uparrow \infty$, $\bar{q} \uparrow 1$.

In the above, $\beta = 0$ corresponds to the ergodic cost.

5 Analysis of the Discounted Problem

We fix $\beta > 0$ and study all time-inhomogeneous discounted NE starting from an arbitrary initial distribution. In view of Proposition 3.2, and the change of variables introduced in Section 3.3, these NE are given by the bounded solutions of the nonlinear dynamical system (2, 8) with $\beta > 0$.

5.1 Linear stability analysis of equilibria

The local behavior of the dynamical system around the origin is described by the spectral properties of the linearized system,

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix} a(t) \\ q(t) \end{pmatrix} = \begin{pmatrix} \beta + 2\sigma^2 & -\kappa \\ 1 & -2\sigma^2 \end{pmatrix} \begin{pmatrix} a(t) \\ q(t) \end{pmatrix}.$$

An analysis of the eigenvalues of this system shows that:

- *Subcritical regime* $\kappa < \kappa_c$: The origin is a saddle point (one positive, one negative eigenvalue).
- *Supercritical regime (A)* $\kappa_c < \kappa < \kappa_c + \beta^2/4$: The origin is unstable (two positive eigenvalues). The other stationary points are saddles.
- *Supercritical regime (B)* $\kappa > \kappa_c + \beta^2/4$: The origin is a spiral source (two complex eigenvalues with positive real part). The other two stationary points remain saddles.

5.2 Global analysis

We rewrite the equations (2, 8) as $(\dot{a}, \dot{q})^\intercal = f(a, q)$, where

$$f(a, q) := \begin{pmatrix} (\beta + 2\sigma^2)a + \mathrm{sign}(a)a^2/2 - \kappa q \\ a - (2\sigma^2 + |a|)q \end{pmatrix}.$$

A direct calculation shows that $\mathrm{div}(f) \equiv \beta$. Moreover, for a given initial condition $(a, q) \in D$, we let $\Phi(t, a, q)$ be the unique solution at time $t \in I(a, q)$. Here $I(a, q) \subset \mathbb{R}$ is the maximal interval where the solution is defined. An *orbit* of the dynamical system (2, 8) is any set given by $\{\Phi(t, a, q) : t \in I(a, q)\}$ for some $(a, q) \in D$.

We start with a result that is repeatedly used in our arguments and which follows from the fact that $\mathrm{div}(f) > 0$.

Lemma 5.1 *If $U \subset D$ is an open bounded set whose boundary is the closure of finitely many orbits, then $U = \emptyset$. In particular, the dynamical system (2, 8) does not have closed orbits and all bounded solutions of this dynamical system converge to one of the stationary points of the system. Similarly, as time goes to minus infinity, all bounded solutions either converge to one of the stationary points or hit the lines $\{q = \pm 1\}$.*

Proof. Toward a contraposition, suppose that $U \neq \emptyset$. Then, using the Gauss' lemma (divergence theorem), we integrate along the ∂U to arrive at

$$0 < \beta\, \mathrm{Leb}(U) = \int_U \mathrm{div} f(x)\, \mathrm{d}x = \int_{\partial U} f(x) \cdot \nu(x)\, \mathrm{d}S(x) = 0,$$

where S is the "surface" measure and $\nu(x)$ denotes the exterior unit normal, which is orthogonal to $f(x)$ as ∂U is an orbit, up to finitely many points. This contradiction implies that there are no bounded closed orbits. Also, by the Poincaré-Bendixson theorem, [18](Theorem II.1.3), all bounded solutions of a planar dynamical system must converge to a stationary point or a bounded closed orbit. Hence, as time goes to plus or minus infinity the bounded solutions either converge to a fixed point or hit $\{q = \pm 1\}$. Moreover, the flow cannot cross $\{q = \pm 1\}$ going forward. □

Note that this lemma rules out homoclinic orbits, and using a symmetry argument, heteroclinic orbits connecting non-trivial fixed points as well. We now study the structure of the Nash equilibria.

Proposition 5.2 (Subcritical regime) *Suppose that $\kappa < \kappa_c$. Then, there exist a strictly increasing function $\mathfrak{n} : [-1, 1] \to \mathbb{R}$ with $\mathfrak{n}(0) = 0$ such that its graph $\{(\mathfrak{n}(q), q) : q \in [-1, 1]\}$ is the stable manifold of the origin of the dynamical system (2, 8). All Nash equilibria are included in the graph of $\mathfrak{n}$. In particular, for any $q \in [-1, 1]$ there exists exactly one discounted Nash equilibrium starting from q, and all Nash equilibria converge to the origin.*

Proof. It suffices to show that the stable manifold $\mathcal{M}_s$ of the origin is a monotone curve in phase space that hits the boundary $\{q = \pm 1\}$ as $t \to -\infty$. Additionally, by symmetry it is enough to provide the proof for the upper boundary. Let $\mathcal{Q}$ denote the region above the nullclines S_a and S_q, as shown in Figure 5. Any trajectory passing through $\mathcal{Q}$ is strictly monotone in the phase space as both components of f are negative on $\mathcal{Q}$, and hence is represented by a function. Moreover, linear analysis around the origin shows that the upper stable manifold $\mathcal{M}_s \cap \{a > 0\}$ lies locally inside $\mathcal{Q}$, and therefore it is locally represented by a strictly monotone graph. If time is reversed, this manifold cannot leave $\mathcal{Q}$ as S_a can only be crossed vertically. We thus conclude that it crosses the upper boundary $\{q = 1\}$ in finite time as there are no fixed points of the dynamical system in $\mathcal{Q}$. The same analysis in $\{a < 0\}$ completes the construction of $\mathfrak{n}$. In summary, there is a strictly increasing function $\mathfrak{n} : [-1, 1] \to \mathbb{R}$ with $\mathfrak{n}(0) = 0$ such that $\mathcal{M}_s = \{(\mathfrak{n}(q), q) : q \in [-1, 1]\}$. Then, for any $q \in [-1, 1]$, the solution of (2, 8) starting at $(\mathfrak{n}(q), q)$ remains in the graph of $\mathfrak{n}$ and converges to the origin as time goes to infinity. Moreover, by Lemma 5.1, all bounded solutions must converge to the origin which is the unique stationary point. Thus, they must be contained the stable manifold $\mathcal{M}_s$ of the origin. As Nash equilibria are precisely the bounded solutions of the dynamical system, we conclude that all of them must be in $\mathcal{M}_s$ and thus in the graph of $\mathfrak{n}$. □

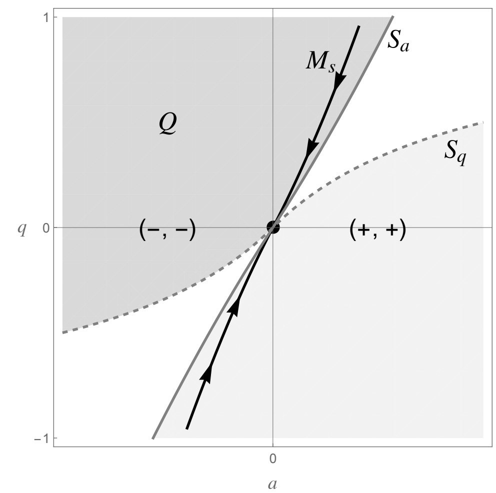

Figure 5: Subcritical discounted system (2, 8).

Recall that in the supercritical regime, we have three stationary points of the dynamical system: $(-\overline{a}, -\overline{q})$, $(0, 0)$, and $(\overline{a}, \overline{q})$. We now study the supercritical regime whose relevant phase diagrams are drawn in Figures 3a and 3b.

Theorem 5.3 (Supercritical regime) *Suppose that $\kappa > \kappa_c$. Then, there is a curve $\mathcal{C} \subset D$ that connects all three stationary points of the dynamical system (2, 8) and hits the boundary $\{q = \pm 1\}$. It is given by the stable manifolds of the non-trivial equilibria joining the origin. All discounted NE are included in $\mathcal{C}$ so that for any $q \in [-1, +1]$ there is at least one NE starting from q. Moreover,*

(A) if $\kappa < \kappa_c + \beta^2/4$, then $\mathcal{C}$ is a strictly monotone curve, so that there is a unique discounted NE starting from q, and when $q \neq 0$ it converges to sign$(q)\overline{q}$.

(B) if $\kappa > \kappa_c + \beta^2/4$, then $\mathcal{C}$ spirals around the origin. In particular, for initial data q close to the origin there exist many NE that spiral around the origin before converging to one of the self-organizing SNE. However, for $|q|$ sufficiently large, there is a unique NE starting from q.

Proof. We first establish the picture depicted in Figure 6. Let $\mathcal{M}_u$ denote the unstable manifold of the positive stationary point $(\overline{a}, \overline{q})$. We claim that $\mathcal{M}_u$ extends to the left as a monotone graph that lies above $(-\overline{a}, -\overline{q})$. We continue by proving this claim. Let $\mathcal{Q}_1$ be the region above both S_a and S_q, let $\mathcal{Q}_2$ be the lens enclosed by S_a and S_q in $\{a < 0\}$, let $\mathcal{Q}_3$ be the area below both S_a and S_q, and finally let $\mathcal{Q}_4$ be the lens enclosed by S_a and S_q in $\{a > 0\}$, as shown in Figure 6. Consider a trajectory starting at a point in $\mathcal{M}_u \cap \{a < \overline{a}\}$ which is close to $(\overline{a}, \overline{q})$. Linear analysis shows that locally this trajectory lies in $\mathcal{Q}_1$, and it stays above $S_q \cap \{a > 0\}$ which can only be intersected horizontally. We claim that once it enters $S_q \cap \{a < 0\}$, it stays in $\mathcal{Q}_1$ by following a monotone graph that lies above $(-\overline{a}, -\overline{q})$. Indeed, while being in $\{-\overline{a} < a < 0\}$, it can only leave $\mathcal{Q}_1$ by entering into $\mathcal{Q}_2$. Suppose this happens. Then, it either stays bounded or unbounded. Suppose it stays bounded. As the origin is unstable and the stationary point $(-\overline{a}, -\overline{q})$ can be reached from $\mathcal{Q}_1$, the trajectory would have to converge to $(\overline{a}, \overline{q})$, contradicting Lemma 5.1. Hence it has to be unbounded by entering into $\mathcal{Q}_3$ through $\mathcal{Q}_2$.

If this happens, by symmetry, we can construct another unbounded trajectory starting from $(-\overline{a}, -\overline{q})$ which enters into $\mathcal{Q}_1$ from $\mathcal{Q}_4$. Then, these two trajectories would intersect, leading to a contradiction. Thus the original trajectory does not enter into $\mathcal{Q}_2$ and stays in $\mathcal{Q}_1$, proving the claim and establishing Figure 6 where the unstable manifolds are drawn by thick lines.

Next we construct $\mathcal{C}$. We claim that the lower stable manifold $\mathcal{M}_s \cap \{q < \overline{q}\}$ of $(\overline{a}, \overline{q})$ extends to the origin. Indeed, if it is unbounded, by symmetry, we conclude that it would cross the upper stable manifold $\tilde{\mathcal{M}}_s \cap \{q > -\overline{q}\}$ of $(-\overline{a}, -\overline{q})$ yielding a contradiction. Hence, it must remain bounded. Moreover, it cannot intersect the unstable manifold emanating from $(-\overline{a}, -\overline{q})$. This prevents it hitting the lower boundary $\{q = -1\}$. Then, by Lemma 5.1 in reverse time, it converges to any of the fixed points of the flow. We analyze all cases separately. As there are no homoclinic orbits, it cannot converge to $(\overline{a}, \overline{q})$. If it converges to $(-\overline{a}, -\overline{q})$ connecting two non-zero fixed points, we can construct a symmetric orbit connecting them in the opposite direction. The union of these heteroclinic orbits is also not possible by Lemma 5.1. Hence, it must converge to the origin.

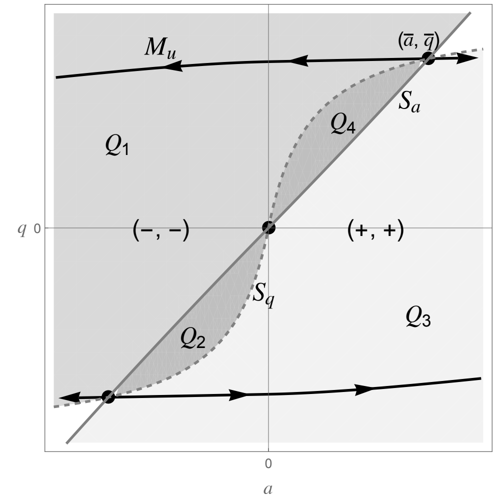

Figure 6: Supercritical discounted system (2, 8).

We have shown that $\mathcal{M}_s$ joins the origin. By symmetry, this implies that the stable manifold $\tilde{\mathcal{M}}_s$ of $(-\overline{a}, -\overline{q})$ joins the origin as well. Moreover, a direct argument shows that the stable manifolds of the non-trivial equilibria extend to the boundary $\{q = \pm 1\}$ as monotone curves. In summary, the union $\mathcal{C}$ of the stable manifolds of the non-trivial fixed points together with the origin is a curve that extends from $\{q = 1\}$ to $\{q = -1\}$ going through all three fixed points. Moreover, an application of Lemma 5.1 shows that all bounded solutions lie on $\mathcal{C}$. To construct a discounted NE starting from $q \in [-1, 1]$, we choose $a^*(q)$ such that $(a^*(q), q) \in \mathcal{C}$. Then, the q-component of the solution starting from this point is a discounted NE starting from q.

In case (A), we claim that $\mathcal{C}$ constructed above is a monotone graph. Indeed, linear analysis around the origin implies that $\mathcal{M}_s$ connecting the origin to $(\overline{a}, \overline{q})$ must enter into $\mathcal{Q}_3$. Linear analysis around $(\overline{a}, \overline{q})$ implies that $\mathcal{M}_s$ enters into $\mathcal{Q}_3$ in reverse time. Hence, for these two parts to connect, all of $\mathcal{M}_s \cap \{0 < q < \bar{q}\}$ must lie in $\mathcal{Q}_3$. Since the components of the vector field of the dynamical system are positive in $\mathcal{Q}_3$, $\mathcal{C} \cap \mathcal{Q}_3$ is a monotone graph. By symmetry we conclude that $\mathcal{C}$ is a monotone graph. Further, the monotonicity of $\mathcal{C}$ implies that $a^*(q)$ is unique. Since all NE starting from q are given as the q-component of a bounded solution that lies in $\mathcal{C}$, we conclude that there is a unique one for every q.

In case (B), by the Hartman-Grobman theorem, [32](Chapter 2.8), the nonlinear dynamical system is topologically conjugate to the linearized system in a neighborhood of the origin. As the linear system spirals around the origin, so does $\mathcal{C}$ of the nonlinear system. Therefore, for small q, there are many points a so that $(a, q) \in \mathcal{C}$, and for each one there is a discounted NE starting from q. When $|q|$ sufficiently large, $a^*(q)$ is unique and so is the NE starting from q. □

5.3 Analysis as a potential game

All Kuramoto mean-field games discussed in this paper are potential, so that the MFG systems are the first-order conditions of an associated mean-field optimal control problem, see [19]. In the special case of finite-state MFGs, this is a direct consequence of the necessary part of Pontryagin's maximum principle, see for example [15], [16] for accounts of this fact.

We continue by describing the optimal control problem associated to the two-state model that we study. We first recall that a feedback control is a measurable function $\alpha : [0, \infty) \to \mathbb{R}$, and the value of the control from state 0 to 1 is $\alpha(t, 0) := \alpha^+(t)$ and from 1 to 0 is $\alpha(t, 1) := \alpha^-(t)$. Let X_t^α be the process corresponding this feedback control, and set $p(t) := \mathbb{P}(X_t^\alpha = 1)$, $q(t) := 2p(t) - 1$ as before. Then, with this interpretation,

$$\mathbb{E}[\alpha(t, X_t^\alpha)^2] = p(t)(\alpha^-(t))^2 + (1 - p(t))(\alpha^+(t))^2 = \frac{1}{2}\,(\alpha(t)^2 - \text{sign}(\alpha(t))\alpha(t)^2 q(t)).$$

The running cost of the mean-field control problem is given by,

$$\begin{aligned}\frac{1}{2}\,\mathbb{E}[\alpha(t, X^\alpha)^2] + \frac{\kappa}{2}\mathbb{E}[\ell(X_t^\alpha, p(t))] &= \frac{1}{4}\,(\alpha(t)^2 - \text{sign}(\alpha(t))\alpha(t)^2 q(t)) + \kappa p(t)(1 - p(t)) \\ &= \frac{1}{4}\,[(\alpha(t)^2 - \text{sign}(\alpha(t))\alpha(t)^2 q(t)) + 1 - q(t)^2].\end{aligned}$$

Then, an equivalent optimal control problem is

$$\inf_{\alpha\in\mathcal{A}} \quad \int_0^\infty e^{-\beta t} \left(\alpha(t)^2 - \operatorname{sign}(\alpha(t))\alpha(t)^2 q(t) - \kappa q(t)\right) \mathrm{d}t, \tag{MFC$_\beta$}$$

where the controlled state process $q(\cdot)$ is a solution of (1) with a replaced with α, and initial condition $q(0) = q$. If α^* is an optimal control, the optimally controlled path $q^*(\cdot)$ is a discounted NE starting from q. Then, the minimizers of the above problem are one of the discounted NE characterized in the previous section. We leverage this connection to obtain several results about the minimizers. The following is an immediate corollary of Proposition 5.2 and Theorem 5.3.

Corollary 5.4 *Suppose that $\kappa < \kappa_c + \beta^2/4$. Then, for every $q \in [-1, 1]$ there is a unique minimizer of (MFC$_\beta$), and the optimally controlled state converges to the uniform distribution.*

Proposition 5.5 *For sufficiently large κ, the optimally controlled state $q^*(\cdot)$ converges to one of the self-organizing SNE.*

Proof. We only have to consider the case of a uniform initial condition, $q(0) = 0$. Then, we simply take a constant control $\alpha \equiv a > 0$ and compute the discounted cost explicitly. A straightforward computation shows that we can achieve a strictly negative cost if

$$\kappa > \kappa_c + \frac{\beta^2}{2} + \beta\sigma^2 =: \widetilde{\kappa}$$

and $a < (\kappa - \widetilde{\kappa})/(\beta + 2\sigma^2)$. This proves that the zero control cannot be optimal for sufficiently large κ, and the minimizer must follow one of the Nash equilibria that converge to the self-organizing SNE. □

6 Analysis of the Ergodic Problem

The critical interaction parameter is $\kappa_c = 4\sigma^4$, and the dynamical system (2, 8) reduces to

$$\dot{a}(t) = 2\sigma^2 a(t) + \frac{1}{2}\operatorname{sign}(a(t))a(t)^2 - \kappa q(t), \tag{10}$$

$$\dot{q}(t) = a(t) - (2\sigma^2 + |a(t)|)q(t). \tag{11}$$

This system is conservative, and admits the following first integral which remains constant along any solution to (10, 11),

$$E(a, q) = \kappa\frac{q^2}{2} + \frac{a^2}{2} - 2\sigma^2 aq - \frac{a^2}{2}\operatorname{sign}(a)q.$$

Proposition 6.1 (Ergodic subcritical regime) *If $\kappa < \kappa_c$, then the uniform SNE is the unique ergodic Nash equilibrium.*

Proof. We already know that the origin is an ergodic NE, and the only stationary one. By Proposition 3.1, any other ergodic NE is the q-component of a periodic trajectory or equivalently a closed orbit. Moreover, linear stability analysis shows that the origin is a saddle point, and any solution (a, q) that goes through the origin satisfies $E(a, q) = E(0, 0) = 0$. Hence, the stable manifold of the origin is given as the graph of the following function,

$$a \mapsto \frac{1}{2\kappa}\left(\operatorname{sign}(a)a^2 + 4\sigma^2 a + a\sqrt{a^2 + 8\sigma^2|a| + 4(\kappa_c - \kappa)}\right).$$

This is a strictly increasing function in phase space that crosses the origin and hits the boundary $\{q = \pm 1\}$.

If there is a closed orbit, it must enclose the origin since it is the only stationary point of this system. Then, this closed orbit would have to cross the stable manifold, creating a contradiction. □

We next analyze the supercritical regime. In this case, a direct calculation implies that the origin is a strict local minimum of the energy E. Since the origin is an isolated stationary point of (10, 11), all trajectories sufficiently close to the origin are closed. Then, the origin becomes a nonlinear center, and we have the following rather unexpected result.

Proposition 6.2 (Ergodic supercritical regime) *If $\kappa > \kappa_c$, then there exist infinitely many time-periodic NE rotating around the uniform.*

6.1 Analysis as a potential game

This sections studies Nash equilibria that arise as first-order conditions of the following ergodic control problem,

$$\inf_{\alpha\in\mathcal{A}} \quad \overline{\lim_{T\uparrow\infty}} \frac{1}{T}\int_0^T \left(\alpha(t)^2 - \operatorname{sign}(\alpha(t))\alpha(t)^2 q(t) - \kappa q(t)^2\right)\, \mathrm{d}t, \tag{MFC$_0$}$$

where $q(\cdot)$ is the solution of (8) with $a = \alpha$ and any initial data. Notice that since $\sigma > 0$, changes in the initial condition do not affect the ergodic cost. Moreover, there are minimizers of the above control problem. As explained in Section 5.3, any optimally controlled path $q^*(\cdot)$ of (MFC$_0$) that is periodic is an ergodic NE.

Theorem 6.3 (Ergodic mean-field control problem) *Fix any initial condition. If $\kappa < \kappa_c$, the zero control is optimal and the associated state converges to the uniform distribution. In the supercritical regime $\kappa > \kappa_c$, however, the zero control is not optimal and the state does not converge to the uniform dsitribution.*

Proof. If $\kappa < \kappa_c$, we have shown in the proof of Proposition 6.1 that any bounded solution (a, q) to (10, 11) converges to the origin. Hence the cost of the ergodic control problem (MFC$_0$) is zero, and we can achieve this cost by taking the zero control. Now let $\kappa > \kappa_c$. Let $\overline{a} > 0$ denote the stationary control corresponding to the positive supercritical SNE, given by (9). We take the constant control $\alpha \equiv \overline{a} > 0$ and directly compute that

$$\overline{a}^2 - \overline{a}^2 q - \kappa\overline{q}^2 < 0 \quad \Longleftrightarrow \quad \kappa > \kappa_c.$$

This implies that we can achieve a strictly negative ergodic cost if $\kappa > \kappa_c$, proving the claim. □

6.2 Convergence of the finite horizon game

Fix $T > 0$ and for a given flow of probabilities $(p(t))_{t\in[0,T]}$ consider the finite horizon problem of the representative agent:

$$v^T(t,x) := \inf_{\alpha\in\mathcal{A}} \mathbb{E}\int_t^T \left(\frac{1}{2}\alpha(u)^2 + \kappa\ell(X_u^\alpha, p(u))\right) \mathrm{d}u,$$

where $X_t^\alpha = x$. Then, the finite horizon Nash equilibria are defined analogously, and following the proof of Proposition 3.1, we conclude that finite horizon Nash equilibria are in one-to-one correspondence with bounded trajectories (a^T, q^T) of the conservative system (10, 11), together with an initial condition for $q^T(0)$ and the terminal condition $a^T(T) = 0$.

In the supercritical case, a direct analysis of the energy E reveals that the stable and unstable manifolds of the non-trivial equilibria join each other, creating a lens as depicted in Figure 7. Inside this lens, any orbit is periodic. Since solutions depend continuously on the initial data, we can compare trajectories starting on the stable manifolds of the non-trivial equilibria to ones close to them to conclude that orbits spend most time near the self-organizing SNE. Further, as periodic orbits inside the lens approach the boundary of the lens, their period goes to infinity. Therefore, we have the following *turnpike property*: For any self-organizing SNE and any initial distribution there exists a sequence of finite horizon NE $q^T(\cdot)$, $T \in \mathbb{N}$, that gets arbitrarily close to the self-organizing one as $T \uparrow \infty$. This property is somehow in analogy with the local stability result proved in [12](Theorem 3.1) for the original Kuramoto mean-field game in the large coupling constant regime. However, for initial conditions in $(-\overline{q}, \overline{q})$, there are many other sequences of finite horizon Nash equilibria that stay bounded away from any SNE as $T \uparrow \infty$ as well.

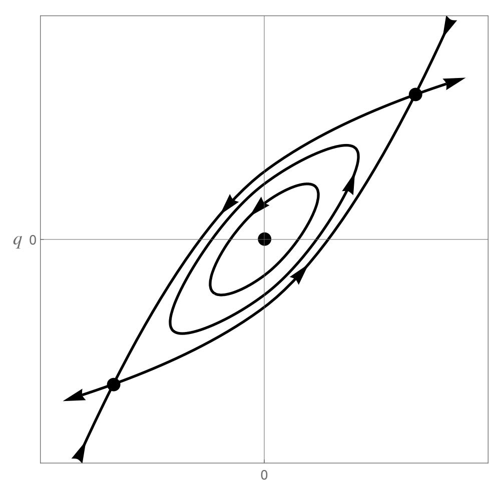

Figure 7: Stable and unstable manifolds of the non-trivial equilibria of (10, 11), and periodic orbits in the supercritical regime $\kappa > \kappa_c$.

7 Conclusion

We study a two-state mean-field game as a tractable model for synchronization obtained as the discretization of the Kuramoto mean-field game introduced and studied in [5], [34]. After characterizing Nash equilibria (NE) as the bounded solutions of a planar dynamical system (1,2), we provide a complete analysis of all cases. In particular, the game exhibits phase transitions as the classical Kuramoto model [27]. While the uniform distribution, representing the incoherent state, is the unique stationary Nash equilibrium (SNE) for small coupling constants, coherent SNE emerge as this interactions get stronger. Moreover, in the supercritical regime, there are many time-inhomogeneous NE starting from the same initial distribution. Surprisingly, there are also periodic NE with the ergodic cost.

References

[1] D. Acemoğlu, G. Como, F. Fagnani, and A. Özdağlar, "Opinion Fluctuations and Disagreement in Social Networks," *Mathematics of Operations Research*, vol. 38, no. 1, pp. 1–27, 2013.

[2] A. Arapostathis, V. S. Borkar, and M. K. Ghosh, *Ergodic control of diffusion processes*. Cambridge University Press, 2012, vol. 143.

[3] D. Bauso, H. Tembine, and T. Başar, "Opinion Dynamics in Social Networks through Mean-Field Games," *SIAM Journal on Control and Optimization*, vol. 54, no. 6, pp. 3225–3257, 2016.

[4] C. Bertucci and A. Cecchin, "Mean field games master equations: From discrete to continuous state space," *arXiv preprint arXiv:2207.03191*, 2022.

[5] R. Carmona, Q. Cormier, and H. M. Soner, "Synchronization in a Kuramoto Mean Field Game," *Communications in Partial Differential Equations*, vol. 48, no. 9, pp. 1214–1244, 2023.

[6] R. Carmona and F. Delarue, *Probabilistic Theory of Mean Field Games with Applications I*. Springer International Publishing, 2018, vol. 83.

[7] R. Carmona and C. V. Graves, "Jet Lag Recovery: Synchronization of Circadian Oscillators as a Mean Field Game," *Dynamic Games and Applications*, vol. 10, no. 1, pp. 79–99, 2020.

[8] R. Carmona and P. Wang, "A Probabilistic Approach to Extended Finite State Mean Field Games," *Mathematics of Operations Research*, vol. 46, no. 2, pp. 471–502, 2021.

[9] C. Castellano, S. Fortunato, and V. Loreto, "Statistical physics of social dynamics," *Reviews of Modern Physics*, vol. 81, no. 2, pp. 591–646, 2009.

[10] A. Cecchin and M. Fischer, "Probabilistic Approach to Finite State Mean Field Games," *Applied Mathematics & Optimization*, vol. 81, no. 2, pp. 253–300, 2020, ISSN: 1432-0606.

[11] A. Cecchin, P. D. Pra, M. Fischer, and G. Pelino, "On the Convergence Problem in Mean Field Games: A Two State Model without Uniqueness," *SIAM Journal on Control and Optimization*, vol. 57, no. 4, pp. 2443–2466, 2019.

[12] A. Cesaroni and M. Cirant, "Stationary equilibria and their stability in a kuramoto mfg with strong interaction," *Communications in Partial Differential Equations*, pp. 1–27, 2024.

[13] A. Cohen and E. Zell, "Analysis of the Finite-State Ergodic Master Equation," *Applied Mathematics & Optimization*, vol. 87, no. 3, p. 40, 2023.

[14] P. Dai Pra, E. Sartori, and M. Tolotti, "Climb on the Bandwagon: Consensus and Periodicity in a Lifetime Utility Model with Strategic Interactions," *Dynamic Games and Applications*, vol. 9, no. 4, pp. 1061–1075, 2019.

[15] D. Gomes, R. M. Velho, and M.-T. Wolfram, "Socio-economic applications of finite state mean field games," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 372, no. 2028, p. 20 130 405, 2014.

[16] D. A. Gomes, J. Mohr, and R. R. Souza, "Continuous Time Finite State Mean Field Games," *Applied Mathematics & Optimization*, vol. 68, no. 1, pp. 99–143, 2013.

[17] O. Guéant, "Existence and Uniqueness Result for Mean Field Games with Congestion Effect on Graphs," *Applied Mathematics & Optimization*, vol. 72, no. 2, pp. 291–303, 2015.

[18] J. K. Hale, *Ordinary Differential Equations*. Dover Publications, 2009.

[19] F. Höfer and H. M. Soner, "Potential games, gradient flows and mean field control," *in preparation*, 2024.

[20] M. Huang, P. E. Caines, and R. P. Malhamé, "An Invariance Principle in Large Population Stochastic Dynamic Games," *Journal of Systems Science and Complexity*, vol. 20, no. 2, pp. 162–172, 2007.

[21] M. Huang, P. E. Caines, and R. P. Malhamé, "Large-Population Cost-Coupled LQG Problems With Nonuniform Agents: Individual-Mass Behavior and Decentralized ε-Nash Equilibria," *IEEE Transactions on Automatic Control*, vol. 52, pp. 1560–1571, 2007, ISSN: 1558-2523.

[22] M. Huang, P. E. Caines, and R. P. Malhamé, "The Nash certainty equivalence principle and McKean-Vlasov systems: An invariance principle and entry adaptation," in *2007 46th IEEE Conference on Decision and Control*, 2007, pp. 121–126.

[23] M. Huang, C. PE, and R. Malhamé, "Individual and mass behaviour in large population stochastic wireless power control problems: Centralized and Nash equilibrium solutions," in *42nd IEEE International Conference on Decision and Control*, vol. 1, 2003, pp. 98–103.

[24] V. N. Kolokoltsov and A. Bensoussan, "Mean-Field-Game Model for Botnet Defense in Cyber-Security," *Applied Mathematics & Optimization*, vol. 74, no. 3, pp. 669–692, 2016.

[25] V. N. Kolokoltsov and O. A. Malafeyev, "Corruption and botnet defense: A mean field game approach," *International Journal of Game Theory*, vol. 47, no. 3, pp. 977–999, 2018.

[26] V. N. Kolokoltsov and O. A. Malafeyev, "Mean-Field-Game Model of Corruption," *Dynamic Games and Applications*, vol. 7, no. 1, pp. 34–47, 2017.

[27] Y. Kuramoto, "Self-entrainment of a population of coupled non-linear oscillators," in *International Symposium on Mathematical Problems in Theoretical Physics*, H. Araki, Ed., ser. Lecture Notes in Physics, Berlin, Heidelberg: Springer, 1975, pp. 420–422.

[28] J.-M. Lasry and P.-L. Lions, "Jeux à champ moyen. I – Le cas stationnaire," *Comptes Rendus Mathematique*, vol. 343, no. 9, pp. 619–625, 2006.

[29] J.-M. Lasry and P.-L. Lions, "Jeux à champ moyen. II – Horizon fini et contrôle optimal," *Comptes Rendus Mathematique*, vol. 343, no. 10, pp. 679–684, 2006.

[30] J.-M. Lasry and P.-L. Lions, "Mean field games," *Japanese Journal of Mathematics*, vol. 2, no. 1, pp. 229–260, 2007.

[31] P.-L. Lions, "Cours au Collège de France," *Available at www.college-de-france.fr*, 2007.

[32] L. Perko, *Differential Equations and Dynamical Systems*. Springer US, 1996, vol. 7.

[33] H. Yin, P. G. Mehta, S. P. Meyn, and U. V. Shanbhag, "Bifurcation analysis of a heterogeneous mean-field oscillator game model," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*, 2011, pp. 3895–3900.

[34] H. Yin, P. G. Mehta, S. P. Meyn, and U. V. Shanbhag, "Synchronization of coupled oscillators is a game," *IEEE Transactions on Automatic Control*, vol. 57, no. 4, pp. 920–935, 2011.