

A mean-field game model for homogenous flocking

Grover, P.; Bakshi, K.; Theodorou, E.

TR2018-032 June 20, 2018

Abstract

Empirically derived continuum models of collective behavior among large populations of dynamic agents are a subject of intense study in several fields, including biology, engineering and finance. We formulate and study a mean-field game model whose behavior mimics an empirically derived non-local homogenous flocking model for agents with intrinsic gradient-type self-propulsion dynamics. The mean-field game framework provides a non-cooperative optimal control description of the behavior of a population of agents in a distributed setting. In this description, each agent's state is driven by optimally controlled dynamics that result in a Nash equilibrium between itself and the population. The optimal control is computed by minimizing a cost that depends only on its own state, and a mean-field term. The agent distribution in phase space evolves under the optimal feedback control policy. We exploit the low-rank perturbative nature of the non-local term in the forward-backward system of equations governing the state and control distributions, and provide a linear stability analysis demonstrating that our model exhibits bifurcations similar to those found in the uncontrolled model. The present work is a step towards developing a set of tools for systematic analysis, and eventually design, of collective behavior of non-cooperative dynamic agents via an inverse modeling approach.

Chaos: An Interdisciplinary Journal of Nonlinear Science

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

A mean-field game model for homogeneous flocking

Piyush Grover,¹ Kaivalya Bakshi,^{1,2} and Evangelos A. Theodorou²

¹*Mitsubishi Electric Research Labs, Cambridge, MA 02139, USA*^{a)}

²*Aerospace Engineering, Georgia Tech, Atlanta, GA 30332, USA*

(Dated: 29 May 2018)

Empirically derived continuum models of collective behavior among large populations of dynamic agents are a subject of intense study in several fields, including biology, engineering and finance. We formulate and study a mean-field game model whose behavior mimics an empirically derived nonlocal homogeneous flocking model for agents with gradient self-propulsion dynamics. The mean-field game framework provides a *non-cooperative optimal control* description of the behavior of a population of agents in a distributed setting. In this description, each agent's state is driven by optimally controlled dynamics that result in a Nash equilibrium between itself and the population. The optimal control is computed by minimizing a cost that depends only on its own state, and a mean-field term. The agent distribution in phase space evolves under the optimal feedback control policy. We exploit the low-rank perturbative nature of the nonlocal term in the forward-backward system of equations governing the state and control distributions, and provide a closed-loop linear stability analysis demonstrating that our model exhibits bifurcations similar to those found in the empirical model. The present work is a step towards developing a set of tools for systematic analysis, and eventually *design*, of collective behavior of non-cooperative dynamic agents via an inverse modeling approach.

While the analysis of emergent behavior in a large population of dynamic agents is a classical topic, the design of desired macroscopic behavior in such systems is a grand engineering challenge. Such systems are often studied using continuum models, involving empirically derived systems of nonlinear partial differential equations that govern the distribution of agents in the phase space. The various terms in these equations represent intrinsic dynamics of the agents, mutual attraction and/or repulsion, and noise. An important class of such models concern flocking, both in nature, and engineering applications such as bio-inspired control of multi-agent robotics, traffic modeling, power-grid synchronization etc. We take a mean-field game approach to derive a control system that mimics the behavior of one such class of models in the setting of non-cooperative agents. A mean-field game is a coupled system of partial differential equations that govern the state and optimal control distributions of a representative agent in a Nash equilibrium with the population. Using a linear stability analysis, we recover phase transitions that have been observed in the corresponding empirical model, as well as find some new ones, as the control penalty is changed.

I. INTRODUCTION

Continuum models of large populations of interacting dynamic agents are popular in mathematical biology¹,

and also have been employed in numerous applications such as multi-agent robotics², finance³ and traffic modeling⁴. The aim of such models is to accurately represent the macroscopic dynamics of the population, and its dependence on parameters. Typically, such models are derived by starting with an empirical dynamical system for a representative agent. This system typically involves the intrinsic dynamics of the agent, a coupling function⁵ describing its interaction with the population, and noise. From this single agent dynamical system, a continuum description is obtained by deriving a macroscopic equation for the distribution of agents in the phase space. We call this class of models *uncontrolled*.

An alternative way of deriving continuum models of collective behavior is via a corresponding variational principle. In this approach, the dynamical system for a representative agent includes its intrinsic dynamics, a control term and noise. The unknown control term is obtained as a solution to an optimization problem. Within this variational (or optimization) framework for large populations, there are multiple classes of modeling strategies⁶. If one takes a centralized global optimization viewpoint, the corresponding problem is that of *mean-field control*, i.e. it is assumed that each agent is being controlled by a central entity whose goal is to optimize a macroscopic cost function⁷ that includes interaction among the population. In a distributed setting, there is no central entity, and the agents can either be *cooperative* or *non-cooperative*. In the former case, each agent chooses its control to optimize a global sum of cost functions of the population.

On the other hand, in the *non-cooperative* mean field setting that we are interested in, each agent optimizes only its individual cost function. This cost function involves coupling with the population solely via a mean-field term. This is the setting of mean-field games (MFG)⁸⁻¹⁰. In this setting, a Hamilton-Jacobi-Bellman (HJB) equation (posed backward in time) characterizes

^{a)}Electronic mail: grover@merl.com

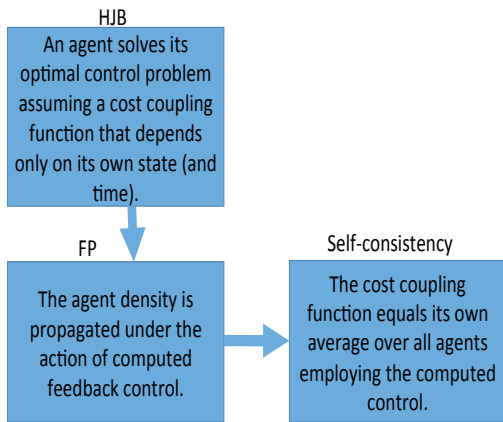


FIG. 1: The MFG framework

the optimal feedback control for a representative agent under the assumption that the (cost) coupling function depends only on its own state, and possibly time. A Fokker-Planck (FP) equation governs the evolution of agent density in phase space. A consistency principle¹⁰ requires that the coupling function used in the agent HJB equation is reproduced as its own average over the continuum of agents. Under fairly general conditions, solutions to MFG model can be shown to possess ϵ -Nash property, i.e., unilateral benefit of any deviation from the computed control policy by a single agent vanishes rapidly as the population becomes large.

The classical (uncontrolled) Cucker-Smale (CS) flocking model¹¹ describes a system of finite population of coupled agents with trivial intrinsic dynamics, moving solely under the influence of an alignment force, and noise. This was followed by several continuum descriptions^{12,13}, and was recently generalized to a continuum model with self-propulsion effects in the homogeneous case¹⁴ (i.e., assuming spatial homogeneity). This latter generalization results in existence of non-zero mean velocity distribution resulting from symmetry breaking, a wide range of ‘disordered’ states consisting of multiple flocks, and other phase transitions.

A MFG model for a continuum of coupled Kuramoto oscillators¹⁵ was described in a seminal work¹⁶ that influences the development in the current paper. Building upon this work, a MFG model for the classical inhomogeneous CS was then proposed¹⁷; the stability analysis was partially addressed. This was followed by a homogeneous flocking MFG model for coupled agents with trivial intrinsic dynamics, along with linear and nonlinear stability analysis¹⁸. Also of interest is an approach¹⁹ where agents apply a gradient descent rather than solve an HJB equation, since the Nash equilibria of the MFG are recovered under certain conditions using this approach.

The contributions of this paper are as follows. We formulate a MFG model for homogeneous flocking of agents driven by self-propulsion and noise. In contrast to the earlier work on homogeneous MFG model with trivial

intrinsic dynamics¹⁸, this model exhibits phase transitions (bifurcations) that mimic those present in the corresponding uncontrolled model¹⁴. We generalize the stability analysis developed in previous MFG models^{10,16,18,20} to agents with gradient nonlinear dynamics, and employ a method used to study reaction-diffusion equations²¹ to derive a semi-analytical stability criterion. Besides qualitatively explaining the phase transition phenomena, quantitative results useful in control design are obtained from the numerical analysis. Decreasing the control control penalty below a threshold causes the zero mean velocity steady state of the MFG model to lose stability via pitchfork bifurcation²². This results in a pair of stable steady states with non-zero mean velocity. If the control is made even cheaper, a new stable regime (nonexistent in the uncontrolled model) emerges for zero mean velocity steady states in the small noise case via a subcritical pitchfork bifurcation.

II. UNCONTROLLED FORMULATION

We briefly review here the uncontrolled formulation from Ref. 14 which provides a homogeneous model for CS flocking with self-propulsion. Consider a population of N agents moving in phase space $((q, p) \in \mathbb{R}^2)$, where each agent is acted upon by a gradient self-propulsion term, a CS coupling force with localization kernel K in position space that aligns the agents’ velocity with the neighbors, and noise. The dynamics for i th agent are

$$dq_i = p_i dt, \\ dp_i = a(p_i)dt + F(q_i, p_i, q_{-i}, p_{-i})dt + \sigma d\omega_i,$$

where $a(p_i) = -\partial_p U(p_i)$, $U(p_i) = \alpha(\frac{p_i^4}{4} - \frac{p_i^2}{2})$, $F(q_i, p_i, q_{-i}, p_{-i}) = \frac{1}{N} \frac{\sum_{j=1}^N K(q_i, q_j)(p_j - p_i)}{\sum_{j=1}^N K(q_i, q_j)}$, $\sigma > 0$ is the noise intensity, $\alpha > 0$ defines the strength of the self-propulsion term, $K(q, q') = K(q', q) \geq 0$, and $K(q, q) = 1$, $q_{-i} = \{q_1, \dots, q_{i-1}, q_{i+1}, \dots\}$, $p_{-i} = \{p_1, \dots, p_{i-1}, p_{i+1}, \dots\}$.

In the continuum limit ($N \rightarrow \infty$), the agent density $f(q, p, t)$ in phase space is governed by

$$\partial_t f + \partial_q(p f) + \partial_p(a(p)f + F[f]f) = \frac{\sigma^2}{2} \partial_{pp} f,$$

where $F[f](q, p, t) = (\bar{p} - p)$ and,

$$\bar{p}(q, t) = \frac{\int \int K(q, q') p f(q', p, t) dq' dp}{\int \int K(q, q') f(q', p, t) dq' dp}.$$

We denote the action of the operator F on a function f by $F[f](\cdot)$. From here onwards, we consider the homogeneous case by dropping dependence on q , and use x to denote the velocity p . The uncontrolled dynamics for the velocity of agent i are

$$dx_i = a(x_i)dt + \frac{1}{N} \sum_{j=1}^N (x_j - x_i)dt + \sigma d\omega_i, \quad (1)$$

with corresponding density evolution

$$\partial_t f = \partial_x(\alpha(x^2 - 1)xf + (x - \bar{x})f) + \frac{\sigma^2}{2}\partial_{xx}f, \quad (2)$$

$$\text{where } \bar{x}(t) = \frac{\int xf(x, t)dx}{\int f(x, t)dx}.$$

A. Fixed Points and Stability Analysis

It is known^{14,23} that fixed points of Eq. (2) are given by

$$f_\infty(x; \mu) = \frac{1}{Z} \exp\left(\frac{-2}{\sigma^2} \left[\alpha \frac{x^4}{4} + (1 - \alpha) \frac{x^2}{2} - \mu x\right]\right), \quad (3)$$

where $\mu \in \mathbb{R}$ is the mean of the distribution, and Z is the normalization factor. For all positive values of parameters (σ, α) , the zero mean velocity solution $f_\infty(\cdot, 0)$ always exists. For a range of parameters, two additional stable non-zero mean velocity solutions are created via a supercritical bifurcation, resulting in loss of stability of the zero mean solution. In Ref. 14, these stability properties were inferred numerically by a Monte-Carlo approach.

We take a different approach, and consider the spectral stability of steady state solutions of Eq. (2). In addition to gaining additional insight into the properties of the uncontrolled system, this also sets the stage for stability analysis of the MFG system in the next section. We consider perturbations of the form $f(x, t) = f_\infty(x)(1 + \epsilon \tilde{f}(x, t))$. Then, the linearization of Eq. (2) is

$$\partial_t \tilde{f}(x, t) = L[\tilde{f}](x, t) = L_{loc}[\tilde{f}](x, t) + L_{nonloc}[\tilde{f}](x, t),$$

where, $\hat{U}(x) = U(x) + x^2/2 - \mu x$,

$$L_{loc}[\tilde{f}](x, t) = -\partial_x \hat{U}(x) \partial_x \tilde{f}(x, t) + (\sigma^2/2) \partial_{xx} \tilde{f}(x, t)$$

is a local linear operator, and

$$L_{nonloc}[\tilde{f}](x, t) = \frac{2}{\sigma^2} \partial_x \hat{U}(x) \int y \tilde{f}(y, t) f_\infty(y) dy$$

is a *nonlocal* linear operator. An operator O is called nonlocal if $O[f](x_1)$ depends on $f(x_2)$ (or the derivatives $\partial_x f(x_2), \partial_{xx} f(x_2)$) for some $x_2 \neq x_1$, and local otherwise. Let $q(x) \equiv \frac{2}{\sigma^2} \partial_x \hat{U}(x)$. Then, $\partial_x f_\infty(x) = -q(x) f_\infty(x)$. We define a Hilbert space $\mathbb{H} = L^2(\mathbb{R}, f_\infty dx)$, i.e., the f_∞ -weighted inner-product space of square-integrable functions on the real line. Then we can write a general form of the full linearized operator as

$$L[\tilde{f}](x, t) = L_{loc}[\tilde{f}](x, t) + s_1(x) \langle g_1(\cdot), \tilde{f}(\cdot, t) \rangle, \quad (4)$$

where $s_1(x) = q(x)$, $g_1(x) = x$ for our case, and the inner product is understood to be $\langle \cdot, \cdot \rangle_{\mathbb{H}}$. We note that L_{loc} is a self-adjoint operator²⁴ on \mathbb{H} which has a non-positive discrete real spectrum of the form $0 = \lambda_1 > \lambda_2 >$

$\lambda_3 \dots$. It has a complete set of orthogonal eigenfunctions $\{\xi_i(x)\}_{i \in \mathbb{N}}$. The first eigenfunction ξ_1 , spanning the kernel of L_{loc} , is a constant function. Following the approach presented in Refs. 21 and 25 for nonlocal eigenvalue problems in reaction-diffusion equations (also see Ref. 26), we consider the following eigenvalue problem

$$\begin{aligned} \lambda w &= L_{loc} w + s_1(x) \langle g_1, w \rangle \\ 0 &= (L_{loc} - \lambda I) w + s_1(x) \langle g_1, w \rangle. \end{aligned} \quad (5)$$

Note that an eigenfunction w of L satisfying $\langle w, g_1 \rangle = 0$ is also an eigenfunction of L_{loc} , i.e. $w = v_i$ for some i with eigenvalue $\lambda = \lambda_i$. We search for eigenfunctions such that $\langle w, g_1 \rangle$ is nonzero. The corresponding eigenvalues are called ‘moving’ eigenvalues in Ref. 25. Multiplying both sides of Eq. (5) with the resolvent $R_\lambda = (L_{loc} - \lambda I)^{-1}$,

$$0 = w + R_\lambda s_1(x) \langle g_1, w \rangle.$$

Taking the inner product of the above equation with g_1 ,

$$0 = \langle g_1, w \rangle + \langle R_\lambda s_1(x), g_1 \rangle \langle g_1, w \rangle. \quad (6)$$

$$\text{For an arbitrary function } z(x), R_\lambda z = \sum_{i=1}^{\infty} \frac{\langle \xi_i, z \rangle}{\lambda_i - \lambda} \xi_i.$$

Evaluating the inner product in Eqs. 6,

$$\langle R_\lambda s_1(x), g_1 \rangle = \langle R_\lambda q(x), x \rangle = \sum_{i=2}^{\infty} \frac{\langle \xi_i, q(x) \rangle}{\lambda_i - \lambda} \langle \xi_i, x \rangle. \quad (7)$$

Using this result in Eq. (6),

$$\langle w, x \rangle \left(1 + \sum_{i=2}^{\infty} \frac{\langle \xi_i, q(x) \rangle}{\lambda_i - \lambda} \langle \xi_i, x \rangle\right) = 0.$$

Hence, either $\langle w, x \rangle = 0$, or $1 + \sum_{i=2}^{\infty} \frac{\langle \xi_i, q(x) \rangle}{\lambda_i - \lambda} \langle \xi_i, x \rangle = 0$.

But we are looking for moving eigenvalues, i.e. w s.t. $\langle w, x \rangle \neq 0$, hence the eigenvalue equation reduces to:

$$h(\lambda) \equiv 1 + \sum_{i=2}^{\infty} \frac{\langle \xi_i, q(x) \rangle}{\lambda_i - \lambda} \langle \xi_i, x \rangle = 0. \quad (8)$$

A sufficient condition for Eq. (8) to have only real roots is that the function $h(\lambda)$ is Herglotz, or equivalently, the product $\langle \xi_i, q(x) \rangle \langle \xi_i, x \rangle$ has the same sign for all i . Using integration by parts on eigenvalue equation for L_{loc} , one can show that $\langle \xi_i, x \rangle = -\frac{\sigma^2}{2\lambda_i} \langle \xi_i, q(x) \rangle$. Thus the Herglotz condition is satisfied since $\lambda_i < 0$ for all $i > 1$.

Numerical Results: We use Chebfun²⁷ to perform all computations. The non-zero mean steady state solutions to Eq. (2) are computed using a simple fixed point iteration for μ . The solutions are shown in Figure 2. The supercritical pitchfork bifurcation that occurs as σ is reduced below critical value $\sigma_c(\alpha)$, is shown for a range of α values. To evaluate $h(\lambda)$ in Eq. (8), we compute the spectrum of L_{loc} for $\mu = 0$. The odd-numbered eigenfunctions are even functions of x , and hence $\langle \xi_{2k+1}, g_1 \rangle = 0$. Therefore, eigenvalues λ_{2k+1} of L_{loc} are also eigenvalues of L ,

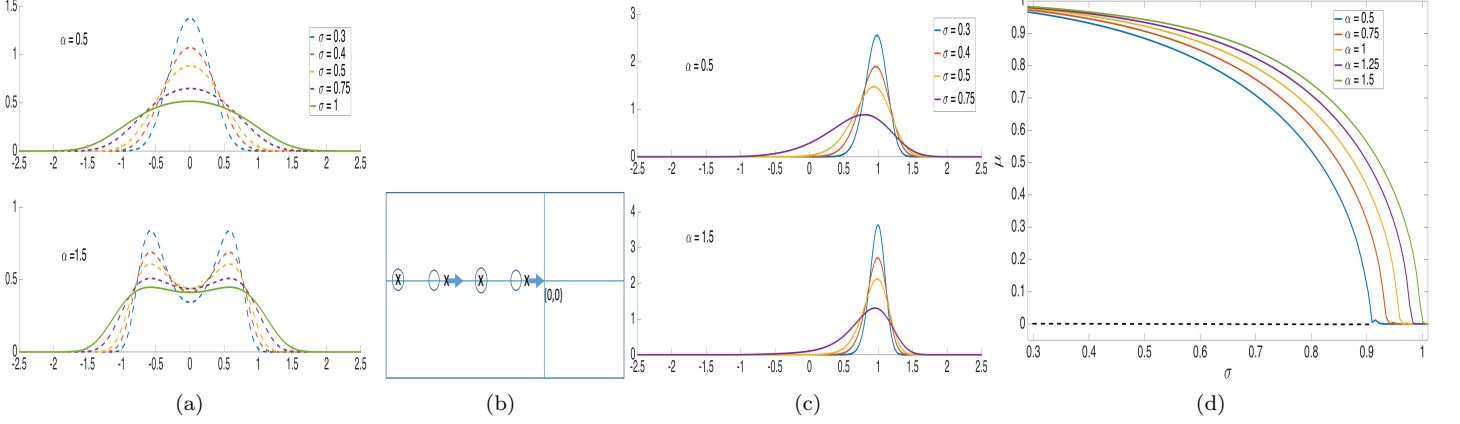


FIG. 2: Uncontrolled system. a) Stable (solid) and unstable (dashed) zero mean steady state solutions for $\alpha = 0.5$ (uni-modal, top) and $\alpha = 1.5$ (bi-modal, bottom). b) Eigenvalues of L_{loc} (o) and L (x) for a typical zero mean case. The first eigenvalue (= 0) is omitted. Notice that alternating eigenvalues are same for both operators. The arrows indicate the direction of motion of the other ('moving') eigenvalues of L as σ is reduced. The rightmost eigenvalue of L reaches 0 at $\sigma = \sigma_c(\alpha)$ with non-zero speed. c) Non-zero mean solutions. d) The $\mu > 0$ branch (solid) bifurcating from $\mu = 0$ solution (dashed) via a supercritical pitchfork bifurcation as σ occurs is reduced below $\sigma_c(\alpha)$.

and the eigenvalues λ_{2k} are moving eigenvalues. We find that at $\sigma = \sigma_c(\alpha)$, $h(0) = 0$. Hence, as σ is decreased below $\sigma_c(\alpha)$, the least stable eigenvalue λ_2 of L_{loc} moves to the positive real axis due to the effect of the nonlocal term, resulting in instability of the zero mean solution.

III. MFG FORMULATION

In this section we describe a MFG formulation for homogeneous equation Eq. (2). The velocity of i th agent evolves via the following equation (compare with Eq (1))

$$dx_i(t) = a(x_i)dt + u_i(t)dt + \sigma dw_i(t), \quad (9)$$

where u_i is the optimal control. Let $F[x_i, x_{-i}](t) \equiv (x_i - \frac{1}{N-1} \sum_{j \neq i} x_j)^2$, and $\beta = \frac{1}{r\sigma^2}$, where $r > 0$ is the control cost or penalty. Here $x_{-i} \equiv \{x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots\}$. Then the i th agent is minimizing the following long time average cost

$$J = \limsup_{T \rightarrow \infty} \frac{1}{T} \left[\int_0^T \beta F[x_i, x_{-i}](t) + \frac{1}{2\sigma^2} u_i(t)^2 \right] dt,$$

that depends on states of all other agents.

To derive the MFG equations (recall Fig. 1), we rewrite the single-agent cost in terms of $\hat{F}(x_i, t)$, the unknown coupling function with dependence on x_i only

$$J = \limsup_{T \rightarrow \infty} \frac{1}{T} \left[\int_0^T \beta \hat{F}(x_i, t) + \frac{1}{2\sigma^2} u_i(t)^2 \right] dt.$$

The resulting single agent HJB equation²⁸ is

$$\begin{aligned} \partial_t v_i(x, t) = & c - \beta \hat{F}(x, t) - a(x) \partial_x v_i(x, t) \\ & + \frac{\sigma^2}{2} (\partial_x v_i(x, t))^2 - \frac{\sigma^2}{2} \partial_{xx} v_i(x, t), \end{aligned} \quad (10)$$

where $v_i(x, t)$ is the single-agent relative value function, c is the minimum average cost, and $u_i(x, t) = -\sigma^2 \partial_x v_i(x, t)$ given in feedback form. Note that the HJB equation is well-posed backward in time. The self-consistency principle yields the expression for \hat{F} in terms of agent density $f(x, t)$ (in the limit $N \rightarrow \infty$):

$$\hat{F}[f](x, t) = \left[\int (x - y) f(y, t) dy \right]^2. \quad (11)$$

Hence, the following set of FP-HJB MFG equations govern the density and value function evolution:

$$\partial_t f(x, t) + \partial_x [(a(x) - \sigma^2 \partial_x v(x, t)) f(x, t)] = \frac{\sigma^2}{2} \partial_{xx} f(x, t), \quad (12)$$

$$\begin{aligned} \partial_t v(x, t) = & c - \beta \hat{F}[f](x, t) - a(x) \partial_x v(x, t) \\ & + \frac{\sigma^2}{2} (\partial_x v(x, t))^2 - \frac{\sigma^2}{2} \partial_{xx} v(x, t). \end{aligned} \quad (13)$$

The unique invariant density satisfying Eq. (12) is

$$f_\infty(x) = \frac{1}{Z} \exp\left(-\frac{2}{\sigma^2} (U(x) + \sigma^2 v_\infty(x))\right). \quad (14)$$

Inserting this expression into Eq. (13), and using the Cole-Hopf transformation²⁹ $\phi(x) = \exp(-v_\infty(x))$, results in the following nonlinear nonlocal eigenvalue problem for $\phi(x)$

$$\begin{aligned} c\phi(x) = & \beta \left\{ x - \int y \exp\left(\frac{-2}{\sigma^2} U(y)\right) \phi^2(y) dy \right\}^2 \phi(x) \\ & - a(x) \phi(x) - \frac{\sigma^2}{2} \partial_{xx} \phi(x), \end{aligned} \quad (15)$$

with the constraint $\int \exp\left(\frac{-2}{\sigma^2} U(y)\right) \phi^2(y) dy = 1$ to ensure normalization of f_∞ . The ground state of this problem yields the desired steady state solutions, with corresponding eigenvalue being the minimum cost c .

A. Stability Analysis

In this section we extend the resolvent based analysis from section II A to the MFG system, and find conditions for *closed-loop* stability of an arbitrary steady state $(f_\infty(x), v_\infty(x))$ to an initial perturbation in density. We consider mass preserving perturbations in density of the form $f(x, t) = f_\infty(x)(1 + \epsilon \tilde{f}(x, t))$, i.e., the initial conditions satisfy $\int f_\infty(x) \tilde{f}(x, 0) dx = 0$. The perturbed value function is taken to be of the form $v(x, t) = v_\infty(x) + \epsilon \tilde{v}(x, t)$. A given steady state is called linearly stable if any perturbation to the density decays to zero under the action of the control, where both the density and control evolution are computed using linearized MFG equations.

Linearization of MFG equations (12,13) yields the non-local system

$$\begin{bmatrix} \partial_t \tilde{f}(x, t) \\ \partial_t \tilde{v}(x, t) \end{bmatrix} = L_{loc}^{FB} \begin{bmatrix} \tilde{f}(x, t) \\ \tilde{v}(x, t) \end{bmatrix}, \quad (16)$$

where $L^{FB} = L_{loc}^{FB} + L_{nonloc}^{FB}$,

$$L_{loc}^{FB} = \begin{bmatrix} L_{loc} & 2L_{loc} \\ 0 & -L_{loc} \end{bmatrix}, L_{nonloc}^{FB} = \begin{bmatrix} 0 & 0 \\ 2\beta s_1 \langle g_1, \cdot \rangle & 0 \end{bmatrix},$$

$s_1(x) = x - \mu$, $g_1(x) = x$, $L_{loc} = -\partial_x(\hat{U})\partial_x + \frac{\sigma^2}{2}\partial_{xx}$, with eigenvalue/eigenfunction pairs denoted by $\{\lambda_i, \xi_i\}$, and $\hat{U}(x) = U(x) + \sigma^2 v_\infty(x)$ in analogy with the definition of L_{loc} in Section II. In addition to the Hilbert space $\mathbb{H} = L^2(\mathbb{R}, f_\infty dx)$ and R_λ as defined earlier, we also consider a subspace $\tilde{\mathbb{H}} = \{f \in \mathbb{H} | \langle f, 1 \rangle = 0\}$.

1. Eigenspectrum of the linearized forward-backward operator

We start off by noting that the characteristic equation of L_{loc}^{FB} is $(L_{loc} - \lambda I)(L_{loc} + \lambda I) = 0$. Hence, its eigenvalues are $\cup_{i \in \mathbb{N}} \{\pm \lambda_i\}$. Now consider the eigenvalue problem for L^{FB} with eigenvalue λ and eigenfunction $[w_f(x) \ w_v(x)]^T$:

$$\lambda \begin{bmatrix} w_f \\ w_v \end{bmatrix} = \begin{bmatrix} L_{loc} w_f + 2L_{loc} w_v \\ 2\beta s_1 \langle g_1, w_f \rangle - L_{loc} w_v \end{bmatrix}. \quad (17)$$

Assuming $\lambda \notin \cup_{i \in \mathbb{N}} \{\pm \lambda_i\}$, R_λ and $R_{-\lambda}$ are well defined. The second equation of Eq. (17) gives

$$w_v = 2\beta R_{-\lambda} s_1 \langle g_1, w_f \rangle.$$

Substituting this expression in the first equation of Eq. (17), and re-arranging,

$$w_f = -4\beta R_\lambda L_{loc} R_{-\lambda} s_1 \langle g_1, w_f \rangle. \quad (18)$$

Taking the inner product of the above equation with g_1 ,

$$\langle g_1, w_f \rangle (1 + 4\beta \langle g_1, R_\lambda L_{loc} R_{-\lambda} s_1 \rangle) = 0.$$

The eigenvalue equation for the $\langle g_1, w_f \rangle \neq 0$ case for moving eigenvalues (as in Section II A) is

$$h(\lambda) \equiv 1 + 4\beta \langle g_1, R_\lambda L_{loc} R_{-\lambda} s_1 \rangle = 0. \quad (19)$$

Using the definition of resolvent in Eq. (19),

$$h(\lambda) = 1 + 4\beta \sum_{i=2}^{\infty} \lambda_i \frac{\langle g_1, \xi_i \rangle \langle s_1, \xi_i \rangle}{\lambda_i^2 - \lambda^2}, \quad (20)$$

and hence,

$$h(\lambda) = 1 + 4\beta \sum_{i=2}^{\infty} \lambda_i \frac{\langle x, \xi_i \rangle^2}{\lambda_i^2 - \lambda^2}. \quad (21)$$

Since Eq. (21) is Herglotz in λ^2 , this implies that the eigenvalues come in pairs, either real or purely imaginary.

Let $\omega \equiv h(0) = 1 + 4\beta \sum_{i=2}^{\infty} \frac{\langle x, \xi_i \rangle^2}{\lambda_i}$.

Lemma III.1. *Consider the eigenvalue equation $h(\lambda) = 0$ for moving eigenvalues.*

- (i) *If $\langle x, \xi_i \rangle \neq 0$ for all $i \geq 2$, then there exists a pair of real roots $\pm \delta_i$ for each $i \geq 2$, such that $\lambda_{i+1} < \delta_i < \lambda_i$.*
- (ii) *Recall that $\lambda_1 = 0$. If $\langle x, \xi_2 \rangle \neq 0$ and $\omega > 0$, there exists a pair of real roots $\pm \delta_1$, such that $\lambda_2 < \delta_1 < 0$.*
- (iii) *If $\langle x, \xi_2 \rangle \neq 0$ and $\omega < 0$, there exists a pair of purely imaginary roots $\pm i\gamma$.*

Proof. (i) Consider the interval $I_i = (\lambda_{i+1}, \lambda_i)$. As $\lambda \rightarrow \lambda_i^-$, $h(\lambda) \rightarrow \infty$, and as $\lambda \rightarrow \lambda_{i+1}^+$, $h(\lambda) \rightarrow -\infty$. It is easy to check that $h(\lambda)$ is monotonic in I_i . By intermediate value theorem, a root δ_i exists in I_i , and by the monotonicity property, it is unique. The result for $-\delta_i$ follows by symmetry.

(ii) Consider the interval $I_1 = (\lambda_2, 0)$. Note that as $\lambda \rightarrow \lambda_2^+$, $h(\lambda) \rightarrow -\infty$, and as $\lambda \rightarrow 0^-$, $h(\lambda) \rightarrow \omega$. Hence, if $\omega > 0$, arguments similar to those in part (i) yield the existence of a real root δ_1 between λ_2 and 0.

(iii) Consider the function $h(i\gamma)$ for real $\gamma > 0$. Clearly, h is monotonic in this interval. Furthermore, as $\gamma \rightarrow \infty$, $h(i\gamma) \rightarrow 0$, and as $\gamma \rightarrow 0^+$, $h(i\gamma) \rightarrow \omega$. By arguments similar to those in part (i), $\omega < 0$ implies that there is a unique root $i\gamma$ of h . \square

2. Contraction analysis of the linearized forward-backward operator

Since the MFG system has a forward-backward nature, spectral information alone is insufficient to derive conclusions about the stability of steady state solutions. A contraction analysis is therefore adopted following Refs. 10 and 16. Consider the linear dynamical system given by Eq. (16), with initial perturbation

in density $f(x, 0) = f_\infty(1 + \epsilon \tilde{f}(x, 0))$. Assuming that $\tilde{v}(x, T) \rightarrow 0$ as $T \rightarrow \infty$, the conditions for existence of a unique solution satisfying this assumption are derived. These conditions also provide a stability criterion. Integrating the \tilde{v} equation in Eq. (16) from t to T ,

$$\begin{aligned} \tilde{v}(x, T) &= e^{-L_{loc}(T-t)} \tilde{v}(x, t) \\ &+ 2\beta \int_t^T e^{-L_{loc}(T-s)} s_1(x) \langle g_1, \tilde{f}(\cdot, s) \rangle ds. \end{aligned}$$

Taking the limit $T \rightarrow \infty$,

$$\tilde{v}(x, t) = -2\beta e^{-L_{loc}t} \int_t^\infty e^{L_{loc}s} s_1(x) \langle g_1(\cdot), \tilde{f}(\cdot, s) \rangle ds. \quad (22)$$

Substituting above equation in the \tilde{f} equation,

$$\begin{aligned} \partial_t \tilde{f}(x, t) &= L_{loc} \tilde{f}(x, t) \\ &- 4\beta L_{loc} e^{-L_{loc}t} \int_t^\infty e^{L_{loc}s} s_1(x) \langle g_1(\cdot), \tilde{f}(\cdot, s) \rangle ds. \end{aligned} \quad (23)$$

Integrating from 0 to t yields the fixed point equation,

$$\tilde{f}(x, t) = e^{L_{loc}t} \tilde{f}(x, 0) + M \tilde{f}(x, t), \quad (24)$$

where the operator M acting on $\tilde{f}(x, t)$ is defined as

$$M \tilde{f}(x, t) = -4\beta e^{L_{loc}t} \int_{r=0}^{r=t} e^{-L_{loc}r} L_{loc} e^{-L_{loc}r} \int_{s=r}^{s=\infty} e^{L_{loc}s} s_1(x) \langle g_1(\cdot), \tilde{f}(\cdot, s) \rangle ds dr. \quad (25)$$

Applying the Laplace transform in time to Eq. (25),

$$\hat{M}(\lambda) = -4\beta R_\lambda L_{loc} R_{-\lambda} s_1 \langle g_1, \cdot \rangle. \quad (26)$$

The operator norm $\|M\|$ is given by

$$\begin{aligned} \|M\| &= \sup_{\lambda \in \mathbb{I}} \sup_{\|\tilde{f}\|=1} \|\hat{M}(\lambda) \tilde{f}\|, \\ &= 4\beta \sup_{\lambda \in \mathbb{I}} \sup_{\|\tilde{f}\|=1} \left\| \sum_{i=2}^\infty \frac{\lambda_i \langle s_1, \xi_i \rangle \langle g_1, \tilde{f} \rangle}{\lambda_i^2 - \lambda^2} \xi_i \right\|, \\ &= 4\beta \sup_{\lambda \in \mathbb{I}} \sup_{\|\tilde{f}\|=1} \sqrt{\sum_{i=2}^\infty \left[\frac{\lambda_i \langle s_1, \xi_i \rangle \langle g_1, \tilde{f} \rangle}{\lambda_i^2 - \lambda^2} \right]^2}, \\ &= 4\beta \|g_1\| \sqrt{\sum_{i=2}^\infty \frac{\langle s_1, \xi_i \rangle^2}{\lambda_i^2}} = 4\beta \|x\| \sqrt{\sum_{i=2}^\infty \frac{\langle x, \xi_i \rangle^2}{\lambda_i^2}}. \end{aligned} \quad (28)$$

Lemma III.2 proved next implies that $\|M\| < 1$ is a sufficient condition for a steady state $(f_\infty(x), v_\infty(x))$ of the nonlinear MFG system Eqs. (12,13) to be linearly stable to density perturbations.

Lemma III.2. *Consider the initial value problem for the linearized system in Eqs. 16, with mass-preserving initial condition $\tilde{f}(x, 0)$ i.e., $\int f_\infty(x) \tilde{f}(x, 0) dx = 0$. If the operator M is a contraction (i.e., $\|M\| < 1$), then the perturbation in density, $\tilde{f}(\cdot, t)$, decays to 0 as $t \rightarrow \infty$. Moreover, $\tilde{v}(\cdot, t)$ also decays to 0 as $t \rightarrow \infty$.*

Proof. If M is a contraction, then we can (formally) invert the Eq. (24), and write the unique solution

$$\begin{aligned} \tilde{f}(x, t) &= (\mathbf{I} - M)^{-1} e^{L_{loc}t} \tilde{f}(x, 0) \\ &= (\mathbf{I} + M + M^2 + \dots) e^{L_{loc}t} \tilde{f}(x, 0). \end{aligned} \quad (29)$$

We note that mass conservation property is equivalent to $\langle \tilde{f}(x, 0), 1 \rangle = 0$, i.e. $\tilde{f}(x, 0) \in \mathbb{H}$. Recall that L_{loc} restricted to \mathbb{H} is a self-adjoint operator with negative eigenvalues $\lambda_i, i = 2, 3, \dots$. Then, $\lim_{t \rightarrow \infty} \|e^{L_{loc}t}\|_{\mathbb{H}} = \lim_{t \rightarrow \infty} e^{\lambda_2 t} = 0$. This proves the decay of $\tilde{f}(\cdot, t)$. The corresponding result for $\tilde{v}(\cdot, t)$ is obtained by inserting the expression for $\tilde{f}(\cdot, t)$ into Eq. (22). \square

Now consider a case where eigenvalue equation in Eq. (18) has a pair of purely imaginary roots $\pm i\gamma (\neq 0)$. Then there is a eigenfunction z_f s.t.

$$\begin{aligned} z_f &= -4\beta R_{i\gamma} L_{loc} R_{-i\gamma} s_1 \langle g_1, z_f \rangle \\ &= \hat{M}(i\gamma) z_f, \end{aligned}$$

by noting Eq. (26). But this implies that norm of \hat{M} is at least 1, hence it is not a contraction. This implies that a necessary condition for M to be a contraction is the absence of non-zero spectra of L^{FB} on the imaginary axis.

B. Numerical Results

Recall that in the MFG problem described by Eqs. (12,13), the representative agent is minimizing a weighted sum of two costs: one penalizes deviation of its velocity from the mean velocity of the agent population, and the other penalizes the control action. In this section, we compute fixed points, and identify phase transitions of this system of equations as the problem parameters are varied. Rather than solving the resulting constrained nonlinear eigenvalue problem 15 directly, we use an iterative algorithm to compute steady state solutions of the MFG system.




FIG. 3: The MFG system with $\sigma = 0.5$. a) Zero mean MFG steady state densities for $\alpha = 0.5$ (top) and $\alpha = 1.5$ (bottom) for various control penalty values. b) (Top) Eigenvalues of L_{loc}^{FB} (o) and L^{FB} (x) for a typical zero mean case. The twin zero eigenvalues of L^{FB} are omitted. The arrows indicate the direction of motion of the ‘moving’ eigenvalues of L^{FB} as r is reduced starting from $r > r_{sup}(\alpha, \sigma)$. Note that the assumption in Lemma III.1(i) is violated in this particular case due to the symmetric nature of the self-propulsion term, and hence, only alternating eigenvalues are actually ‘moving’. The pair of eigenvalues of L^{FB} closest to imaginary axis, $\pm\delta_1$, reaches 0 at $r = r_{sup}$, and moves up/down the imaginary axis for $r < r_{sup}$. (Bottom) The $\mu > 0$ branch (solid) bifurcating from $\mu = 0$ solution (dashed) via a supercritical pitchfork bifurcation as r is reduced below r_{sup} . c) Non-zero mean MFG steady state densities on the supercritical branch.




FIG. 4: The MFG system with $\sigma = 0.5$. a) The norm of operator M for zero mean steady state as control penalty is varied, for various α . b). The bifurcation diagram for $\alpha = 1.5$, showing supercritical and subcritical (inset) bifurcations. Only the $\mu > 0$ non-zero mean branches are shown. c) Non-zero mean MFG steady state densities on the subcritical branch for $\alpha = 1.5$.

We note that the coupling term $\hat{F}[f](x, t)$ evaluated at any steady state density f_∞ is $\hat{F}[f_\infty](x) = (x - \mu)^2$, where $\mu = \int y f_\infty(y) dy$. Again using Cole-Hopf transformation $\phi(x) = \exp(-v(x))$ on the HJB equation leads to a **linear** eigenvalue problem in ϕ :

$$c\phi(x) = \mathcal{L}[\phi](x), \quad (30)$$

where $\mathcal{L}[\phi] = \beta(x - \mu)^2\phi(x) - a(x)\partial_x\phi(x) - \frac{\sigma^2}{2}\partial_{xx}\phi(x)$. We solve Eq. (30) iteratively along with Eq. (14) to find zero mean MFG steady states (f_∞, v_∞) for a range of r , keeping σ and α fixed (See Fig. 3). These solutions are stable (i.e., $\|M\| < 1$) for large r , implying that when control is expensive, the agents use minimal control action. The resulting steady state distribution is bi-modal due to dominance of the self-propulsion force, and dispersion via noise.

These zero mean solutions lose stability (i.e., $\|M\| > 1$) via a supercritical bifurcation as r is reduced below a critical value r_{sup} . The Eq. (21) for moving eigenvalues

of L^{FB} has a double zero root at $r = r_{sup}$, and a pair of purely imaginary roots emerges as r is reduced below r_{sup} . This implies that the pair of symmetric eigenvalues of L^{FB} closest to the imaginary axis reaches 0 at the critical parameter, and then moves up/down the imaginary axis. The stable non-zero mean MFG steady state solutions on the supercritical branch are computed by combining fixed point iteration in μ with a continuation step. This bifurcation provides a MFG interpretation to the pitchfork bifurcation observed in the uncontrolled system, i.e., cheaper control makes it economical to compensate for noise. Hence, the agents apply larger control action to flock together (and reduce the cost of deviation from the population mean), resulting in symmetry breaking non-zero mean solutions.

When noise strength σ is fixed below a critical value, the zero mean solution branch undergoes a *subcritical* bifurcation as control penalty r is further reduced, i.e. at $r = r_{sub} < r_{sup}$ (See Fig. 4). The corresponding non-zero

mean solutions were computed using bisection method. This bifurcation is not seen in the uncontrolled system. For instance, when $(\sigma = 0.5, \alpha = 1.5)$, it results in creation of uni-modal stable zero mean solutions in the case of cheap control, $r < r_{sub}$, as compared to the bi-modal stable zero mean solution that exist for expensive control, $r > r_{sup}$. Hence, we conclude that for $r < r_{sub}$, the control is cheap enough to counteract the intrinsic dynamics, and make zero mean uni-modal solution stable.

IV. CONCLUSIONS

We have presented a MFG formulation for homogeneous flocking of agents with gradient nonlinearity in their intrinsic dynamics. We have employed tools from theory of reaction-diffusion equations, and exploited the low rank nature of the nonlocal coupling term to study the linear stability of the MFG equations. The explicit formulae for verifying the stability of steady state solutions of the nonlocal forward-backward MFG system require relatively simple numerical computation of spectra of the local self-adjoint Fokker-Planck operators. The MFG system shows rich nonlinear behavior, such as supercritical and subcritical pitchfork bifurcations that result in wide range of collective behaviors, some of which are not present in the uncontrolled model.

Much of the analysis in the current work can be generalized to higher dimensional state space for homogeneous flocking with self-propulsion, similar in spirit to the generalization¹⁴ of one-dimensional *uncontrolled* flocking model. Furthermore, the abstract results presented in this work apply to models other than homogeneous flocking, e.g. nonlocally coupled agents with arbitrary first order gradient dynamics. Extension to non-homogeneous flocking would be a natural next step; the resulting second-order dynamics could require more sophisticated tools³⁰ for stability analysis. Implementation of the MFG control laws in an engineered large population system requires the control to be provided in a causal form. Algorithms that can learn the MFG laws can be used to convert the control laws obtained by solving the FP-HJB equations into an implementable form³¹.

The use of bifurcation and singularity theory to develop bio-inspired control and decision making algorithms for multi-agent systems has been explored recently^{32–34}. Our work adds to the toolbox for systematic analysis of collective behavior of non-cooperative dynamic agents via an inverse modeling approach. The qualitative and quantitative insight provided by the stability analysis can be exploited in *mechanism design*, i.e., design of penalties or incentives to drive the population to asymptotic states with desirable characteristics. We believe that a systematic study of bifurcations in MFG models can lead to progress in tackling the grand challenge of designing or manipulating collective behavior of a large population of non-cooperative dynamic agents.

ACKNOWLEDGMENTS

We wish to thank the anonymous reviewers for their constructive comments, especially regarding Lemma III.1 and its proof.

- ¹H. Resat, L. Petzold, and M. F. Pettigrew, “Kinetic modeling of biological systems,” (Springer, 2009) pp. 311–335.
- ²G. Beni, “From swarm intelligence to swarm robotics,” in *International Workshop on Swarm Robotics* (Springer, 2004) pp. 1–9.
- ³E. Bonabeau, “Agent-based modeling: Methods and techniques for simulating human systems,” *Proceedings of the National Academy of Sciences* **99**, 7280–7287 (2002).
- ⁴G. Whitham, “On kinematic waves ii. a theory of traffic flow on long crowded roads,” in *Proc. R. Soc. Lond. A*, Vol. 229 (The Royal Society, 1955) pp. 317–345.
- ⁵T. Stankovski, T. Pereira, P. V. McClintock, and A. Stefanovska, “Coupling functions: universal insights into dynamical interaction mechanisms,” *Reviews of Modern Physics* **89**, 045001 (2017).
- ⁶M. Nourian, P. E. Caines, R. P. Malhame, and M. Huang, “Nash, social and centralized solutions to consensus problems via mean field control theory,” *IEEE Transactions on Automatic Control* **58** (2013).
- ⁷K. Elamvazhuthi and P. Grover, “Optimal transport over nonlinear systems via infinitesimal generators on graphs,” arXiv preprint arXiv:1612.01193 (2016).
- ⁸P. E. Caines, “Mean field games,” *Encyclopedia of Systems and Control*, 706–712 (2015).
- ⁹J.-M. Lasry and P.-L. Lions, “Mean field games,” *Japanese journal of mathematics* **2**, 229–260 (2007).
- ¹⁰M. 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* **52**, 1560–1571 (2007).
- ¹¹F. Cucker and S. Smale, “Emergent behavior in flocks,” *IEEE Transactions on automatic control* **52**, 852–862 (2007).
- ¹²S.-Y. Ha, J.-G. Liu, *et al.*, “A simple proof of the cucker-smale flocking dynamics and mean-field limit,” *Communications in Mathematical Sciences* **7**, 297–325 (2009).
- ¹³Y.-P. Choi, S.-Y. Ha, and Z. Li, “Emergent dynamics of the Cucker–Smale flocking model and its variants,” in *Active Particles, Volume 1* (Springer, 2017) pp. 299–331.
- ¹⁴A. B. Barbaro, J. A. Cañizo, J. A. Carrillo, and P. Degond, “Phase transitions in a kinetic flocking model of Cucker–Smale type,” *Multiscale Modeling & Simulation* **14**, 1063–1088 (2016).
- ¹⁵Y. Kuramoto, “Self-entrainment of a population of coupled nonlinear oscillators,” in *International symposium on mathematical problems in theoretical physics* (Springer, 1975) pp. 420–422.
- ¹⁶H. Yin, P. G. Mehta, S. P. Meyn, and U. V. Shanbhag, “Synchronization of coupled oscillators is a game,” *IEEE Transactions on Automatic Control* **57**, 920–935 (2012).
- ¹⁷M. Nourian, P. E. Caines, and R. P. Malhamé, “Mean field analysis of controlled cucker-smale type flocking: Linear analysis and perturbation equations,” *IFAC Proceedings Volumes* **44**, 4471–4476 (2011).
- ¹⁸M. Nourian, P. E. Caines, and R. P. Malhame, “A mean field game synthesis of initial mean consensus problems: A continuum approach for non-gaussian behavior,” *IEEE Transactions on Automatic Control* **59**, 449–455 (2014).
- ¹⁹P. Degond, J.-G. Liu, and C. Ringhofer, “Large-scale dynamics of mean-field games driven by local nash equilibria,” *Journal of Nonlinear Science* **24**, 93–115 (2014).
- ²⁰O. Guéant, “A reference case for mean field games models,” *Journal de mathématiques pures et appliquées* **92**, 276–294 (2009).
- ²¹T. J. Anastasio, A. K. Barreiro, and J. C. Bronski, “A geometric method for eigenvalue problems with low-rank perturbations,” *Royal Society open science* **4**, 170390 (2017).
- ²²S. Wiggins, *Introduction to Applied Nonlinear Dynamical Sys-*

- tems and Chaos*, Texts in Applied Mathematics Science, Vol. 2 (Springer-Verlag, Berlin, 1990).
- ²³J. Tugaut, “Phase transitions of McKean–Vlasov processes in double-wells landscape,” *Stochastics An International Journal of Probability and Stochastic Processes* **86**, 257–284 (2014).
- ²⁴G. A. Pavliotis, *Stochastic processes and applications: Diffusion processes, the Fokker-Planck and Langevin equations*, Vol. 60 (Springer, 2014).
- ²⁵P. Freitas, “Stability of stationary solutions for a scalar non-local reaction-diffusion equation,” *The Quarterly Journal of Mechanics and Applied Mathematics* **48**, 557–582 (1995).
- ²⁶T. Kapitula and K. Promislow, “Stability indices for constrained self-adjoint operators,” *Proceedings of the American Mathematical Society* **140**, 865–880 (2012).
- ²⁷T. A. Driscoll, N. Hale, and L. N. Trefethen, “Chebfun guide,” (2014).
- ²⁸V. S. Borkar, “Ergodic control of diffusion processes,” in *Proceedings ICM* (2006).
- ²⁹E. Todorov, “Eigenfunction approximation methods for linearly-solvable optimal control problems,” in *Adaptive Dynamic Programming and Reinforcement Learning, 2009. ADPRL’09. IEEE Symposium on* (IEEE, 2009) pp. 161–168.
- ³⁰J. Alexander, R. Gardner, and C. Jones, “A topological invariant arising in the stability analysis of travelling waves,” *J. reine angew. Math* **410**, 143 (1990).
- ³¹P. Cardaliaguet and S. Hadikhanloo, “Learning in mean field games: The fictitious play,” *ESAIM: Control, Optimisation and Calculus of Variations* **23**, 569–591 (2017).
- ³²N. E. Leonard, “Multi-agent system dynamics: Bifurcation and behavior of animal groups,” *Annual Reviews in Control* **38**, 171–183 (2014).
- ³³V. Srivastava and N. E. Leonard, “Bio-inspired decision-making and control: From honeybees and neurons to network design,” in *American Control Conference (ACC), 2017* (IEEE, 2017) pp. 2026–2039.
- ³⁴R. Gray, A. Franci, V. Srivastava, and N. E. Leonard, “Honeybee-inspired dynamics for multi-agent decision-making,” arXiv preprint arXiv:1503.08526 (2015).