2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

The Geometry of Cyclical Social Trends

Bernard Chazelle Kritkorn Karntikoon Jakob Nogler

***Abstract*— We investigate the emergence of periodic behavior in opinion dynamics and its underlying geometry. For this, we use a bounded-confidence model with contrarian agents in a convolution social network. This means that agents adapt their opinions by interacting with their neighbors in a time-varying social network. Being contrarian, the agents are kept from reaching consensus. This is the key feature that allows the emergence of cyclical trends. We show that the systems either converge to nonconsensual equilibrium or are attracted to periodic or quasi-periodic orbits. We bound the dimension of the attractors and the period of cyclical trends. We exhibit instances where each orbit is dense and uniformly distributed within its attractor. We also investigate the case of randomly changing social networks.**

I. INTRODUCTION

Much of the work in the area of opinion dynamics has focused on consensus and polarization [1], [2]. Typical questions include: How do agents come to agree or disagree? How do exogenous forces drive them to consensus? How long does it take for opinion formation to settle? Largely left out of the discussion has been the emergence of *cyclical trends*. A question worth examining is whether the process conceals deeper mathematical structure. The purpose of this work is to show that it is, indeed, the case.

Our main result is a proof that adding a simple contrarian rule to the classic bounded-confidence model suffices to produce quasi-periodic trajectories. The model is a slight variant of the classic *Hegselmann-Krause (HK)* framework: a finite collection of agents hold opinions on several topics, which they update at discrete time steps by consulting their neighbors in a (time-varying) social network. The modification is the addition of a simple repulsive force field that keep agents away from tight consensus. The idea is partly inspired by swarming dynamics, e.g. birds refrain from flocking too closely. Likewise, near-consensus on a large enough scale tends to induce contrarian reactions among agents [3], [4]. Some political scientists have pointed to contrarianism as one of the reasons for the closeness of some national elections [5], [6].

Based on computer simulation, we found that it is not specific distributions of initial opinions that produce oscillations but, rather, the recurrence of certain symmetries in the networks. We prove that the condition is sufficient (though its necessity is still open). Moreover, we show that contrarian opinions tend to orbit toward an attractor whose dimensionality is *independent* of the number of opinions held by a single agent. These attracting sets are typically Minkowski sums of ellipses.

Our inquiry builds on the pioneering work of French [7], DeGroot [8], Friedkin & Johnsen [9], and Deffuant et al. [10]. The model we use is a minor modification of the *bounded-confidence model* model [11], [12]. A *HK system* consists of n agents, each one represented by a point in $\mathbb{R}^d$. The d coordinates for each agent i represent their current opinions on d different topics: thus, d is the dimension of the opinion space. At any (discrete) time, each agent i moves to the mass center of the agents within a fixed distance r_i, which represents its radius of influence. This step is repeated ad infinitum. Formally, the agents are positioned at $x_1(t), \ldots, x_n(t) \in \mathbb{R}^d$ at time t and for any $t = 0, 1, 2, \ldots,$

$$x_i(t+1) = \frac{1}{|\mathcal{N}_i(t)|} \sum_{j \in \mathcal{N}_i(t)} x_j(t), \tag{1}$$

with $\mathcal{N}_i(t) = \left\{ 1 \le j \le n \,:\, \left\| x_i(t) - x_j(t) \right\|_2 \le r_i \right\}$.

Interpreting each $\mathcal{N}_i(t)$ as the set of neighbors of agent i defines the *social network* G_t at time t. In the special case where all the radii of influence are equal ($r_i = R$), convergence into fixed-point clusters occurs within a polynomial number of steps [13]–[15]. Computer simulation suggests that the same remains true even when the radii differ but a proof has remained elusive. In this work, we present a model where cyclical trends in opinion change arise. This is accomplished by using a vertex-transitive graph as social network (specifically a Cayley graph), which stipulates that agents cannot be distinguished by their local environment. Before defining the model formally in the next section, we summarize our main findings.

- Undirected networks always drive the agents to nonconsensual convergence, i.e., to fixed points at which they "agree to disagree." For their behavior to become periodic or quasi-periodic, the social networks need to be directed. We prove that such systems either converge or are attracted to periodic or quasi-periodic orbits. We give precise formulas for the orbits.
- We investigate the geometry of the attractors. We bound the *rotation number*, which indicates the speed at which (quasi)-periodic opinions undergo a full cycle. We exhibit instances where each limiting orbit forms a set that is dense and, in fact, uniformly distributed on its attractor.
- We explore the case of social networks changing randomly at each step. We prove the surprising result that the dimension of the attractor can *decrease* because of the randomization. This is a rare case where adding entropy to a system can reduce its dimensionality.

This work was supported in part by NSF grant CCF-2006125.
Princeton University, chazelle@cs.princeton.edu.
Princeton University, kritkorn@cs.princeton.edu.
ETH Zurich, Switzerland, jnogler@ethz.ch.

The dynamics of contrarian views has been studied before [3]–[6], [16]–[18] but, to our knowledge, not for the purpose of explaining cyclical trends.

II. Contrarian Opinion Dynamics

The social network, with each vertex representing an agent, is modeled as a time-dependent Cayley graph $G_{C_t} = (V, E_t)$ over an abelian group. Since every finite abelian group is isomorphic to a direct sum of cyclic groups, the vertex set V can be written as $(\mathbb{Z}/n_1\mathbb{Z}) \oplus \cdots \oplus (\mathbb{Z}/n_m\mathbb{Z})$. For simplicity, we set $n_i = n$, allowing us to represent $V = (\mathbb{Z}/n\mathbb{Z})^m$ as a vector space. The number of agents is denoted by $N = |V| = n^m$. The edge set $E_t = \bigcup\{(v, v + c) \mid v \in V, c \in C_t\}$ at time t is determined by the *convolution set* $C_t \subseteq (\mathbb{Z}/n\mathbb{Z})^m$.

We model the opinion $x_v(t)$ of agent v at time t as a point in $\mathbb{R}^d$, with the initial opinion $x_v(0)$ abbreviated to x_v. The agent's subsequent opinions evolve based on a weighted average of their own previous opinion (allowing for self-confidence) and the opinions of neighboring agents. In the spirit of *HK* systems, we define the dynamics as follows: for $t = 0, 1, \ldots,$

$$x_v(t+1) = p x_v(t) + \frac{1-p}{|C_t|} \sum_{w \in v + C_t} x_w(t). \qquad (2)$$

Here, p represents a "self-confidence" weight such that $1/N < p < 1$. This choice implies that the agent places more trust in their own opinion than in that of at least one other agent, without being excessively confident. Because of the presence of p, we may assume that C_t does not contain the origin $\mathbf{0}$.

If we view each $x_v(t)$ as a row vector in $\mathbb{R}^d$, the update (2) specifies an N-by-N stochastic matrix F_{C_t}. Let $x(t)$ denote the N-by-d matrix whose rows are the N agent positions $x_v(t)$, for $v \in V$. We have $x(t+1) = F_{C_t} x(t)$. The matrix F_{C_t} may not be symmetric but it is always doubly-stochastic. This means that the mass center $\mathbf{1}^\top x(t)/N$ is time-invariant. Since the dynamics itself is translation-invariant, we are free to move the mass center to the origin, which we do by assuming $\mathbf{1}^\top x = \mathbf{0}^\top$, where x denotes $x(0)$.

Obviously, some initial conditions are uninteresting: for example, $x = \mathbf{0}$. For this reason, we choose x randomly; specifically, each x_v is picked iid from the d−dimensional normal distribution $\mathcal{N}(\mathbf{0}, 1)$. In the following, we use the phrase "with high probability," to refer to an event occurring with probability at least $1 - \varepsilon$, for any fixed $\varepsilon > 0$. Once we have picked the matrix x randomly, we place the mass center of the agents at the origin by subtracting its displacement from the origin: $x \leftarrow x - \frac{1}{N}\mathbf{1}\mathbf{1}^\top x$.

The agents will be attracted to the origin to form a single-point cluster of consensus in the limit. Responding to their contrarian nature, the agents will restore mutual differences by boosting the own opinions. For that reason we consider the scaled dynamics: $y(0) = x$ and, for $t \geq 0$,

$$y(t+1) = \xi_t F_{C_t} y(t), \qquad (3)$$

where ξ_t is chosen so that the diameter of the system remains roughly constant. As scaling leaves the salient topological and geometric properties of the dynamics unchanged, the precise definition of ξ_t can vary to fit analytical (or even visual) needs.

A. Preliminaries

To analyze the model, it is sufficient to focus on a single connected component of G_{C_t}. As C_t spans the vector space V if and only if G_{C_t} is strongly connected,[1] we assume that C_t spans the vector space V, which implies that $|C_t| \geq m$.

For clarity, throughout the remainder of this section, we write C_t as C, omitting the subscript t. The presence of the weight $p > 0$ in (2) ensures that the diagonal of F_C is positive. Together with the strong connnectivity assumption, this makes the matrix F_C primitive, meaning that $F_C^k > 0$, for some $k > 0$. By the Perron-Frobenius theorem [20], all the eigenvalues of F_C lie strictly inside the unit circle in $\mathbb{C}$, except for the dominant eigenvalue 1, which has multiplicity 1. For any $u, v \in V$, we write $\psi_u^v = \omega^{\langle u,v \rangle}$, where $\omega := e^{2\pi i/n}$. We define the vector $\psi^v = (\psi_u^v \mid u \in V)$ and easily verify that $\{\psi^v \mid v \in V\}$ forms an orthogonal eigenbasis for F_C. The eigenvalue λ_v corresponding to ψ^v satisfies

$$\lambda_v \psi_u^v = p\psi_u^v + \frac{1-p}{|C|} \sum_{w \in u + C} \psi_w^v = p\psi_u^v + \frac{1-p}{|C|} \Big(\sum_{h \in C} \psi_h^v \Big) \psi_u^v .$$

We conclude:

Lemma 2.1. *Each $v \in V$ corresponds to a distinct eigenvector ψ^v, which together form an orthogonal basis for $\mathbb{C}^N$. The corresponding eigenvalue is given by*

$$\lambda_v = p + \frac{1-p}{|C|} \sum_{h \in C} \omega^{\langle v,h \rangle}.$$

We define $\lambda = \max_{v \in V}\{|\lambda_v| < 1\}$ and denote by $W = \{v \in V : |\lambda_v| = \lambda\}$ the set of subdominant eigenvectors. The argument of λ_v plays a key role in our discussion, so we define θ_v such that $\lambda_v = |\lambda_v|\omega^{\theta_v}$, with $\theta_v \in (-n/2, n/2]$ (we will prove in (6), $\lambda_v \neq 0$ for $v \in W$, so θ_v is well defined).

B. The evolution of opinions

We begin with the case of a fixed convolution set $C_t = C$. The initial position of the agents is expressed in eigenspace as $x = \frac{1}{N}\sum_{v \in V} \psi^v (\psi^v)^{\mathrm{H}} x$. Let z_v denote the row vector $(\psi^v)^{\mathrm{H}} x = \sum_{u \in V} \omega^{-\langle v,u \rangle} x_u$. Because $(\psi^v)^{\mathrm{H}} x = \mathbf{1}^\top x = \mathbf{0}^\top$, for $v = \mathbf{0} \in V$,

$$x(t) = \frac{1}{N} \sum_{v \in V \setminus \{\mathbf{0}\}} \lambda_v^t \psi^v z_v. \qquad (4)$$

Lemma 2.2. *With high probability, for all $v \neq \mathbf{0}$,*

$$\Omega\big(\sqrt{1/N}\big) = \|z_v\|_2 = O\big(\sqrt{dN \log dN}\big).$$

Proof. Let $a = (a_u)_{u \in V}$ be the first column of the matrix x. For each $u \in V$, by the initialization of the system, $a_u = \zeta_u - \delta$, where $\zeta_u \sim \mathcal{N}(0, 1)$ and $\delta = \frac{1}{N}\mathbf{1}^\top \zeta$. Given $v \neq \mathbf{0}$, ψ^v is orthogonal to $\psi^{\mathbf{0}} = \mathbf{1}$; hence $(\psi^v)^{\mathrm{H}} a = (\psi^v)^{\mathrm{H}}(\zeta - \delta\mathbf{1}) = (\psi^v)^{\mathrm{H}} \zeta$. Since the random vector ζ is unbiased and $|\omega^{-\langle v,u \rangle}| = 1$, it follows that $\mathbf{var}\big[(\psi^v)^{\mathrm{H}} a\big] = \sum_{u \in V} \mathbf{var}\zeta_u = N$. Thus, the first coordinate $z_{v,1}$ of z_v is of the form $a + ib$, where a and b are sampled (not independently) from $\mathcal{N}(0, \sigma_1^2)$ and $\mathcal{N}(0, \sigma_2^2)$, respectively, such that $\sigma_1^2 + \sigma_2^2 = N$. Thus,

[1] Proof is provided in the full version of the paper [19].

$|z_{v,1}| \le \delta$ with probability at most $2\delta/\sqrt{\pi N}$. Conversely, by the inequality $\operatorname{erfc}(z) \le e^{-z^2}$ for $z > 0$, we find that $|z_{v,1}| = O(\sqrt{N \log(dN/\varepsilon)}\,)$, with probability at least $1 - \varepsilon/dN$, for any $0 < \varepsilon < 1$; hence $\|z_v\|_2 = O(\sqrt{dN \log(dN/\varepsilon)}\,)$, with probability at least $1 - \varepsilon/N$. Setting $\delta = \varepsilon\sqrt{\pi/4N}$ and using a union bound completes the proof. □

We upscale by setting $\xi_t = 1/\lambda$; hence $y(t+1) = y(t)/\lambda$.

Theorem 2.3. *Let a_h and b_h be the row vectors whose u-th coordinates ($u \in V$) are $\cos(2\pi\langle h, u\rangle/n)$ and $\sin(2\pi\langle h, u\rangle/n)$, respectively. With high probability, for each $v \in V$, the agent v is attracted to the trajectory of $y^*_v(t)$, where*

$$y^*_v(t) = \frac{1}{N} \sum_{h \in W} \left(\cos \frac{2\pi(t\theta_h + \langle h, v\rangle)}{n}, \sin \frac{2\pi(t\theta_h + \langle h, v\rangle)}{n}\right) \begin{pmatrix} a_h \\ b_h \end{pmatrix} x. \quad (5)$$

Let $\mu := \max\{|\lambda_v|/\lambda < 1\}$ be the third largest (upscaled) eigenvalue, measured in distinct moduli. The error of the approximation decays exponentially fast as a function of μ:

$$\frac{\|y^*_v(t) - y_v(t)\|_F}{\|y_v(t)\|_F} = O\big(\mu^t N^2 \sqrt{d \log dN}\,\big).$$

Proof. Since the eigenvalues sum up to $\mathrm{tr} F_C = pN$ and 1 has multiplicity 1, we have $pN \le 1 + (N-1)\lambda$; hence, by $p > 1/N$,

$$\lambda \ge \frac{pN - 1}{N - 1} > 0. \quad (6)$$

Writing $\mu_v = \lambda_v/\lambda$ and $\mu = \max\{|\mu_v| < 1\}$, we have $|\mu_v| = 1$ for $v \in W$; recall that $W = \{v \in V : |\lambda_v| = \lambda\}$. By (4), it follows that

$$y(t) = \frac{1}{N} \sum_{v \in W} \mu_v^t \psi^v z_v + \eta(t), \quad (7)$$

where, by Lem. 2.2, with high probability,

$$\begin{aligned} \|\eta(t)\|_F &= \Big\| \frac{1}{N} \sum_{v \in V \setminus (W \cup \{\mathbf{0}\})} \mu_v^t \psi^v z_v \Big\|_F \\ &\le \frac{1}{N} \sum_{v \in V \setminus (W \cup \{\mathbf{0}\})} \mu^t \|\psi^v\|_2 \, \|z_v\|_2 \\ &= O\big(\mu^t N \sqrt{d \log dN}\,\big). \end{aligned}$$

The lower bound of the lemma implies that, for any $v \in W$,

$$\begin{aligned} \Big\| \sum_{v \in W} \mu_v^t \psi^v z_v \Big\|_F^2 &= \mathrm{tr}\Big(\sum_{v \in W} \mu_v^t \psi^v z_v\Big)^{\mathrm{H}} \Big(\sum_{v \in W} \mu_v^t \psi^v z_v\Big) \\ &= \mathrm{tr}\Big\{\sum_{v \in W} z_v^{\mathrm{H}} (\psi^v)^{\mathrm{H}} \psi^v z_v\Big\} \\ &= N \cdot \mathrm{tr}\Big\{\sum_{v \in W} z_v^{\mathrm{H}} z_v\Big\} \\ &= N \sum_{v \in W} \|z_v\|_2^2 \ge \Omega(1). \end{aligned}$$

For large enough $t = \Omega\big(\log(dN)/\log(1/\mu)\big)$, the sum in (7) dominates $\eta(t)$ with high probability, while the latter decays exponentially fast. Thus the dynamics $y(t)$ is asymptotically equivalent to $y^*(t) = \frac{1}{N}\sum_{v \in W} \mu_v^t \psi^v z_v$. Recall that $\lambda_v = |\lambda_v|\omega^{\theta_v}$; since, for $v \in W$, $\mu_v = \lambda_v/\lambda$ has modulus 1, it is equal to ω^{θ_v}. This implies that $y^*_v(t) = \frac{1}{N}\sum_{h \in W}\sum_{u \in V} \omega^{t\theta_h + \langle h, v-u\rangle} x_u$. Because $y^*_v(t)$ is real, we can ignore the imaginary part when expanding the expression above, which completes the proof. □

C. Geometric investigations

The trajectory $y^*_v(t)$ is called the *limiting orbit*.[2] Thm. 2.3 indicates that, with high probability, every orbit is attracted to its limiting form at an exponential rate, so we may focus on the latter. Given the initial placement x of the agents, all the limiting orbits lie in the set $\mathbb{S}$, expressed in parametric form by

$$\mathbb{S} = \frac{1}{N} \sum_{h \in W} \big\{(a_h x)\cos X_h + (b_h x)\sin X_h\big\}. \quad (8)$$

Recall that $a_h x$ and $b_h x$ are row vectors in $\mathbb{R}^d$. The attractor $\mathbb{S}$ is the Minkowski sum of a number of ellipses. We examine the geometric structure $\mathbb{S}$ and explain how the limiting orbits embed into it. To do that, we break up the sum (5) into three parts. Given $h \in W$, we know that $\lambda_h \neq 0$ by (6), so there remain the following cases for the subdominant eigenvalues:

- *real* $\lambda_h > 0$: the contribution to the sum is $c_v x$, where c_v is the row vector
$$c_v := \frac{1}{N} \sum_{h \in W:\, \theta_h = 0} \Big\{ a_h \cos \frac{2\pi\langle h, v\rangle}{n} + b_h \sin \frac{2\pi\langle h, v\rangle}{n} \Big\}. \quad (9)$$
- *real* $\lambda_h < 0$: the contribution is $(-1)^t d_v x$, where, likewise, d_v is the row vector
$$d_v := \frac{1}{N} \sum_{h \in W:\, \theta_h = n/2} \Big\{ a_h \cos \frac{2\pi\langle h, v\rangle}{n} + b_h \sin \frac{2\pi\langle h, v\rangle}{n} \Big\}. \quad (10)$$
- *nonreal* λ_h: we can assume that $\theta_h > 0$ since the conjugate eigenvalue $\bar{\lambda}_h = \lambda_{-h}$ is also present in W. The contribution of an eigenvalue is the same as that of its conjugate since $a_h = a_{-h}$ and $b_h = -b_{-h}$. So the contribution of a given $\theta > 0$ is equal to $e_{v,\theta} x$, where
$$\begin{aligned} e_{v,\theta} &:= \frac{2}{N} \sum_{h \in W:\, \theta_h = \theta} \Big\{ a_h \cos \frac{2\pi(t\theta + \langle h, v\rangle)}{n} + b_h \sin \frac{2\pi(t\theta + \langle h, v\rangle)}{n} \Big\} \\ &= a_{v,\theta} \cos \frac{2\pi\theta t}{n} + b_{v,\theta} \sin \frac{2\pi\theta t}{n}, \end{aligned}$$
and[3]
$$\begin{pmatrix} a_{v,\theta} \\ b_{v,\theta} \end{pmatrix} = \frac{2}{N} \sum_{h \in W:\, \theta_h = \theta} R\left(\frac{-2\pi\langle h, v\rangle}{n}\right) \begin{pmatrix} a_h \\ b_h \end{pmatrix}. \quad (11)$$

Putting all three contributions together, we find

$$y^*_v(t) = c_v x + (-1)^t d_v x + \sum_{\theta \in \vartheta} \Big\{ a_{v,\theta} \cos \frac{2\pi\theta t}{n} + b_{v,\theta} \sin \frac{2\pi\theta t}{n} \Big\} x, \quad (12)$$

where ϑ is the set of distinct $\theta_h > 0$ for $h \in W$ and all other quantities are defined in (9, 10, 11). See Fig. 1 for an illustration of a doubly-elliptical orbit around its torus-like attractor.

[2] The phase space of the dynamical system is $\mathbb{R}^{dN}$, but by abuse of notation we use the word "orbit" to refer the trajectory of a single agent, which lies in $\mathbb{R}^d$.

[3] $R(\alpha) = \begin{pmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{pmatrix}$.

Fig. 1. Two orbits of a single agent around its attractor.

1) Generic elliptical attraction: We prove that, for almost all values of the self-confidence weight p, the set W generates either a single real eigenvalue or two complex conjugate ones. By (12), this shows that almost every limiting orbit is either a single fixed point or a subset of an ellipse in $\mathbb{R}^d$.

Theorem 2.4. *There exists a set Λ of at most $\binom{N}{2}$ reals such that the set W is associated with either a single real eigenvalue or two complex conjugate ones, for any $p \in (1/N, 1) \setminus \Lambda$.*

The system is called *regular* if $p \in (1/N, 1) \setminus \Lambda$. For such a system, either (i) $\vartheta = \{\theta\}$ and $c_v = d_v = \mathbf{0}$, or (ii) $\vartheta = \emptyset$ and exactly one of c_v or d_v equals $\mathbf{0}$. In other words, by (12), we have three cases depending on the subdominant eigenvalues:

$$y_v^*(t) = \begin{cases} c_v x & : \textit{real positive} \\ (-1)^t d_v x & : \textit{real negative} \\ \big(a_{v,\theta} \cos \frac{2\pi\theta t}{n} + b_{v,\theta} \sin \frac{2\pi\theta t}{n} \big) x & : \textit{conjugate pair.} \end{cases} \tag{13}$$

Lemma 2.5. *Consider a triangle abc and let $e = pc + (1-p)a$ and $f = pc + (1-p)b$. Let O be the origin and assume that the segments Oe and Of are of the same length (Fig. 2); then the identity $|a|^2 - |b|^2 = \frac{2p}{1-p}(b-a) \cdot c$ holds.*

Proof. Let $d := \frac{1}{2}(e+f)$ be the midpoint of ef. The segment Od lies on the perpendicular bisector of ef, so it is orthogonal to ef; hence to ab. Thus, $d \cdot (b-a) = 0$. Since $d = \frac{1}{2}(2pc + (1-p)a + (1-p)b)$, the lemma follows from $(2pc + (1-p)(a+b)) \cdot (b-a) = 0$. □

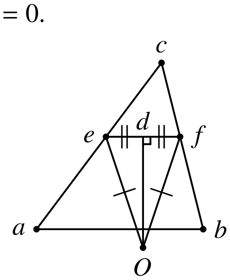

Fig. 2. A triangle identity.

Proof of Theorem 2.4. Choose two distinct $u, v \in W$. Applying Lem. 2.5 in the complex plane, we set: $a = \frac{1}{|C|} \sum_{h \in C} \omega^{\langle u,h \rangle}$; $b = \frac{1}{|C|} \sum_{h \in C} \omega^{\langle v,h \rangle}$; and $c = 1$; thus $e = \lambda_u$ and $f = \lambda_v$, which implies that the segments Oe and Of are of the same length. Abusing notation by treating a, b, c as both vectors and complex numbers, we have $(b-a) \cdot c = \Re(b-a)$; therefore,

$$\big(2\Re(b-a) + |a|^2 - |b|^2\big) p = |a|^2 - |b|^2.$$

1) If $2\Re(b-a) + |a|^2 - |b|^2 = 0$, then $|a| = |b|$, which in turn implies that $\Re(b-a) = 0$; hence $a = \bar{b}$ and $\lambda_u = \bar{\lambda}_v$.
2) If $2\Re(b-a) + |a|^2 - |b|^2 \neq 0$, then p is unique: $p = p_{u,v}$.

We form Λ by including all numbers $p_{u,v}$, with $u, v \in W$. □

2) The case of cycle convolutions: It is useful to consider the case of a single cycle: $m = 1$. For convenience, we momentarily assume that n is prime and that $\sum_{h \in C} h \neq 0 \pmod{n}$; both assumptions are dropped in subsequent sections.

Lemma 2.6. *Each eigenvalue λ_v is simple.*

Proof. Because n is prime, the cyclotomic polynomial for ω is $\Phi(z) = z^{n-1} + z^{n-2} + \cdots + z + 1$. It is the minimal polynomial for ω, which is unique. Note that $\langle v, h \rangle = vh$, since $m = 1$. Given $v \in V$, we define the polynomial $g_v(z) = \sum_{h \in C} z^{vh}$ in the quotient ring of rational polynomials $\mathbb{Q}[z]/(z^n - 1)$. Sorting the summands by degree modulo n, we have $g_v(z) = \sum_{k=0}^{n-1} q_{v,k} z^k$, for nonnegative integers $q_{v,k}$, where $\sum_k q_{v,k} = |C|$. If $\lambda_v = \lambda_u$, for some $u \in V$, then, by Lem. 2.1, $g_v(\omega) = g_u(\omega)$; hence Φ divides $g_v - g_u$. Because the latter is of degree at most $n-1$, it is either identically zero or equal to Φ up to a rational factor $r \neq 0$. In the second case,

$$(q_{v,n-1} - q_{u,n-1}) z^{n-1} + \cdots + (q_{v,1} - q_{u,1}) z + q_{v,0} - q_{u,0} = r\Phi.$$

This implies that $q_{v,k} - q_{u,k} = r \neq 0$, for all $0 \leq k < n$, which contradicts the fact that $\sum_k q_{v,k} = \sum_k q_{u,k} = |C|$; therefore, $g_v = g_u$.

1) If $v = 0$, then $g_v(z) = |C|$; hence $g_u(z) = |C|$ and $u = 0$, i.e., $v = u$.
2) If $v \neq 0$, then let $S_v = \{\omega^{vh} \mid h \in C\}$. Because $\mathbb{Z}/n\mathbb{Z}$ is a field, the $|C|$ roots of unity in S_v are distinct; hence $q_{v,k} \in \{0, 1\}$. It follows that $S_v = S_u$ and $|S_v| = |S_u| = |C|$; therefore, for some permutation σ of order $|C|$, we have $vh = u\sigma(h)$, for all $h \in C$. Summing up both sides over $h \in C$ gives us $v \sum_{h \in C} h = u \sum_{h \in C} h \pmod{n}$; hence $v = u$, since $\sum_{h \in C} h \neq 0 \pmod{n}$. □

By (13), the limiting orbit is of the form $y_v^*(t) = c_v x$ or $y_v^*(t) = (-1)^t d_v x$ if the subdominant eigenvalue is real. Otherwise, the orbit of an agent approaches a single ellipse in $\mathbb{R}^d$: for some $\theta > 0$, $y_v^*(t) = \Big(a_{v,\theta} \cos \frac{2\pi\theta t}{n} + b_{v,\theta} \sin \frac{2\pi\theta t}{n} \Big) x$.

3) Opinion velocities: Assume that the system is regular, so W is associated with either a single real eigenvalue or two complex conjugate ones. If $\vartheta = \emptyset$, by (12), every agent converges to a fixed point of the attractor $\mathbb{S}$ or its limiting orbit has a period of 2. The other case $\vartheta = \{\theta\}$ is more interesting. The agent approaches its limiting orbit, which is periodic or quasi-periodic. The *rotation number*, $\alpha := \theta/n$, is the (average) fraction of a whole rotation covered in a single step. It measures the speed at which the agent cycles around its orbit. It is possible to prove a lower bound on that speed.[4]

Theorem 2.7. *The rotation number α of a regular system satisfies $\alpha \geq \frac{1-p}{n} \left(\frac{1}{2N} \right)^N$.*

[4]Its upper bound is $1/2$.

The proof of Theorem 2.7 is given in [19]. Our next result formalizes the intuitive fact that self-confidence slows down motion. Self-assured agents are reluctant to change opinions.

Theorem 2.8. *The rotation number of a regular system cannot increase with p.*

Proof. We must have $|\vartheta| = 1$. Let λ_h be (an) eigenvalue corresponding to the unique angle in ϑ; recall that $0 < \theta_h < \pi/2$. As we replace p by $p' > p$, we use the prime sign with all relevant quantities post-substitution. Thus, the subdominant eigenvalue for p' associated with ϑ' is denoted by λ'_v; again, we assume that $|\vartheta'| = 1$. Note that v might not necessarily be equal to h; hence the case analysis:

- $v = h$: Using the same notation for complex numbers and the points in the plane they represent (Fig. 3), we see that λ'_h lies in (the relative interior of) the segment $1\lambda_h$; hence $\theta'_h < \theta_h$.
- $v \neq h$: We prove that, as illustrated in Fig. 3, all three conditions $|\lambda_h| > |\lambda_v|$, $|\lambda'_h| < |\lambda'_v|$, and $\theta_h < \theta'_v \leq \pi/2$, cannot hold at the same time, which will establish the theorem. If we increase q continuously from p to p', $\theta_h(q)$ decreases continuously (we use the argument q to denote the fact that θ_h corresponds to the eigenvalue defined with p replaced by q). Since, at the end of that motion, $|\lambda_h(q)| < |\lambda_v(q)|$, by continuity we have $p_o < p'$, where $p_o = \min\{q > p \,:\, |\lambda_h(q)| = |\lambda_v(q)|\}$. To simplify the notation, we repurpose the use of the prime superscript to refer to p_o (eg, $p' = p_o$). So, we now have $|\lambda'_h| = |\lambda'_v|$ and $\theta_h < \theta'_v < \theta_v \leq \pi/2$. It follows that (i) the point λ_v lies in the pie slice of radius $|\lambda_h|$ running counterclockwise from λ_h to $-|\lambda_h|$ on the real axis. Also, because $|\lambda'_h| = |\lambda'_v|$ and $|\lambda_h| > |\lambda_v|$, setting $c = 1$ as before in Lem. 2.5 shows that (ii) $\Re(\lambda_v) > \Re(\lambda_h)$.[5] Putting (i, ii) together shows that $\theta_h \geq \pi/4$ (as shown in Fig. 3). Consequently, the slope of the segment $\lambda_h\lambda_v$ is negative. Since that segment is parallel to $\lambda'_h\lambda'_v$, the perpendicular bisector of the latter has positive slope. Since that bisector is above λ'_v and $\Im(\lambda'_v) \geq 0$, this implies that 0 and λ'_h are on opposite sides of that bisector; hence $|\lambda'_v| < |\lambda'_h|$, which is a contradiction. □

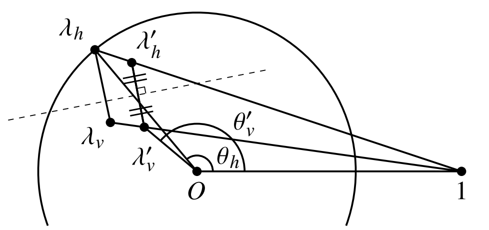

Fig. 3. Why self-confidence slows down the dynamics: proof by contradiction.

D. Equidistributed orbits

The attractor $\mathbb{S}$ is the Minkowski sum of a number of ellipses bounded by $|W|$. An agent orbits around an ellipse as it gets attracted to it exponentially fast. In a regular system with $\vartheta \neq \emptyset$, its limiting orbit is periodic if the unique angle θ_h of ϑ is rational; it is quasi-periodic otherwise. In fact, it then forms a dense subset of the ellipse. By (12), this follows from Weyl's ergodicity principle [21], which states that the set $\{\alpha t \pmod 1), \,|\, t \geq 0\}$ is uniformly distributed in $[0, 1)$, for any irrational α.

Dropping the regularity requirement may produce more exotic dynamics. We exhibit instances where a limiting orbit will not only be dense over the entire attracting set but, in fact, uniformly distributed. In other words, an agent will approach *every possible opinion* with equal frequency. This will occur when this property holds:[6]

Assumption 2.9. *The numbers in $\vartheta \cup \{1\}$ are linearly independent over the rationals.*

We explain this phenomenon next. Order the angles of ϑ arbitrarily and define the vector $\alpha = (\alpha_1, \ldots, \alpha_s) \in \left[0, \frac{1}{2}\right]^s$, where $s = |\vartheta|$ and $\alpha_j = \theta/\pi$ for the j-th angle $\theta \in \vartheta$. We may assume that $c_v = d_v = \mathbf{0}$ in (12) since these cases are rotationally trivial. By Assumption 2.9, $\mathbf{0}$ is the only integer vector whose dot product with α is an integer. We use the standard notation $\|\alpha\|_{\mathbb{R}/\mathbb{Z}} = \max_{k\leq s} \min_{a\in\mathbb{Z}} |\alpha_k - a|$. By Kronecker's approximation theorem [22], for any $\beta \in [0, 1]^s$ and any $\varepsilon > 0$, there exists $q \in \mathbb{Z}$ such that $\|q\alpha - \beta\|_{\mathbb{R}/\mathbb{Z}} \leq \varepsilon$. It follows directly that, with high probability, any limiting orbit is dense over the attractor $\mathbb{S}$. We prove the stronger result:

Theorem 2.10. *Under Assumption 2.9, any limiting orbit is uniformly distributed over the attractor $\mathbb{S}$.*

We mention that, in general, Assumption 2.9 might be difficult to verify analytically. Empirically, however, density is fairly obvious to ascertain numerically with suitable visual evidence (Fig. 4). We define the discrepancy $D(S_t)$ of $S_t = (p_1, \ldots, p_t)$, with $p_i \in \mathbb{R}^s$, as

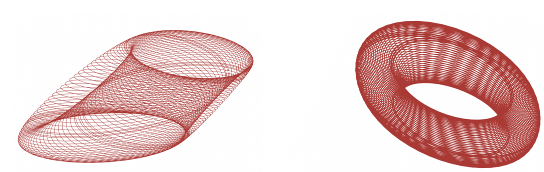

Fig. 4. Two examples where an agent approaches every point on its attractor with equal frequency. In each case, the curve traces the orbit of the agent.

$$D(S_t) = \sup_{B\in J} \left| \frac{A(B;t)}{t} - \mu_s(B) \right|,$$

where μ_s is the s-dimensional Lebesgue measure and $A(B;t) = |\{i \mid p_i \in B\}|$ and J is the set of s-dimensional boxes of the form $\prod_{i=1}^{s} [a_i, b_i) \subset [0, 1]^s$. The infinite sequence S_∞ is said to be *uniformly distributed* if $D(S_t)$ tends to 0, as t goes to infinity.

[5]The keen-eyed observer will notice that in the lemma we must plug in $(p_o - p)/(1 - p)$ instead of p.

[6]The coordinates of $a = (a_1, \ldots, a_k)$ are linearly independent over the rationals if $\mathbf{0}$ is the only rational vector normal to a.

Lemma 2.11. (Erdős–Turán–Koksma [21], page 116). *For any integer $L > 0$,*

$$D(S_t) \le 2s^2 3^{s+1}\Big(\frac{1}{L} + \sum_{0<\|\ell\|_\infty \le L} \frac{1}{r(\ell)}\Big|\frac{1}{t}\sum_{k=1}^{t} e^{2\pi i\langle \ell, p_k\rangle}\Big|\Big),$$

where $r(\ell) := \prod_{j=1}^{s} \max\{1, |\ell_j|\}$ and $\ell = (\ell_1, \ldots, \ell_s) \in \mathbb{Z}^s$.

Proof of Theorem 2.10. We form the sequence $p_1, \ldots, p_t \in [0,1)^s$ such that $p_k = k\alpha \pmod 1$; where each coordinate of $k\alpha$ is replaced by its fractional part. By Lem. 2.11, its box discrepancy satisfies

$$D(S_t) \le 2s^2 3^{s+1}\Big(\frac{1}{L} + \sum_{0<\|\ell\|_\infty \le L} \frac{1}{r(\ell)}\Big|\frac{1}{t}\sum_{k=1}^{t} e^{2\pi i\langle \ell, k\alpha\rangle}\Big|\Big). \tag{14}$$

By Assumption 2.9, $\mathbf{0}$ is the only integer vector whose dot product with α is an integer; hence $\gamma_\ell := e^{2\pi i\langle \ell,\alpha\rangle} \neq 1$, for any $\ell \neq \mathbf{0}$. It follows that

$$\Big|\sum_{k=1}^{t} e^{2\pi i\langle \ell, k\alpha\rangle}\Big| = \Big|\sum_{k=1}^{t} \gamma_\ell^k\Big| = \Big|\frac{\gamma_\ell - \gamma_\ell^{t+1}}{1-\gamma_\ell}\Big| \le \frac{2}{|1-\gamma_\ell|}.$$

By (14), for any $\delta > 0$,

$$D(S_t) \le 2s^2 3^{s+1}\Big(\frac{1}{L} + \frac{1}{t}\sum_{0<\|\ell\|_\infty \le L} \frac{2}{|1-\gamma_\ell|}\Big) \le \delta,$$

for $L = \Big\lceil \frac{4s^2 3^{s+1}}{\delta} \Big\rceil$ and $t \ge (8/\delta)s^2 3^{s+1} \sum_{0<\|\ell\|_\infty \le L} |1-\gamma_\ell|^{-1}$. □

III. Examples

We illustrate the range of contrarian opinion dynamics by considering a few specific examples.

1) Fixed-point attractor: Set $m = 2$ and $C = \{(1,0),(0,1),(-1,0),(0,-1)\}$. By Lem. 2.1, for any $v = (v_1, v_2) \in V$,

$$\lambda_v = p + \frac{1-p}{2}\Big(\cos\frac{2\pi v_1}{n} + \cos\frac{2\pi v_2}{n}\Big).$$

The eigenvalues are real and $\lambda = \max_{v\in V}\{|\lambda_v| < 1\} = p + \frac{1}{2}(1-p)(1+\cos 2\pi/n)$. For any $h \in C$, $\lambda_h = \lambda$ and $\theta_h = 0$; hence $C \subseteq W$. A simple examination shows that, in fact, $W = C$. By (9, 12), given $j \in [d]$,[7]

$$y_v^*(t)_j = A_j \cos\frac{2\pi(v_1+\alpha_j)}{n} + B_j \cos\frac{2\pi(v_2+\beta_j)}{n},$$

where $A_j, B_j, \alpha_j, \beta_j$ do not depend on v but only on the initial position x. This produces a 2D surface in $\mathbb{R}^d$ formed by the Minkowski sum of two ellipses centered at the origin (Fig. 5).

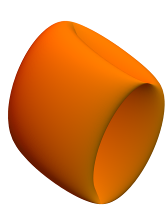

Fig. 5. The attractor on which each agent converges to a fixed point.

[7]As usual, $[d]$ denotes $\{1, \ldots, d\}$.

2) Periodic and quasi-periodic orbits: Set $m = 2$ and $C = \{(1,0),(0,1)\}$. By Lem. 2.1, for any $v \in V$, $\lambda_v = p + \frac{1-p}{2}(\omega^{v_1} + \omega^{v_2})$; hence $\lambda = \max_{v\in V}\{|\lambda_v| < 1\} = \frac{1}{2}\big|1+p+(1-p)\omega\big|$ and $W = \{(1,0),(0,1),(-1,0),(0,-1)\}$. Specifically, λ_v is equal to $\frac{1}{2}\big(1+p+(1-p)\omega\big)$, for $v \in \{(1,0),(0,1)\}$, and to its conjugate, for $v \in \{(-1,0),(0,-1)\}$. By (11, 12), we have $\vartheta = \{\theta\}$, where

$$\theta = \Big(\frac{n}{2\pi}\Big)\arctan\Big(\frac{(1-p)\sin 2\pi/n}{1+p+(1-p)\cos 2\pi/n}\Big), \text{ and}$$

$$y_v^*(t) = \Big(a_{v,\theta}\cos\frac{2\pi\theta t}{n} + b_{v,\theta}\sin\frac{2\pi\theta t}{n}\Big)x.$$

Fix a coordinate $j \in [d]$; we find that

$$y_v^*(t)_j = A_j \cos\frac{2\pi(\theta t + v_1 + \alpha_j)}{n} + B_j \cos\frac{2\pi(\theta t + v_2 + \beta_j)}{n},$$

for suitable reals $A_j, B_j, \alpha_j, \beta_j$ that depend on the initial position x but not on v. This again produces a two-dimensional attracting subset of $\mathbb{R}^d$ formed by the Minkowski sum of two ellipses. In the case of Fig. 6, the attractor is a torus pinched along two curves. The main difference from the previous case comes from the limit behavior of the agents. They are not attracted to a fixed point but, rather, to a surface. With high probability, the orbits are asymptotically periodic if θ is rational, and quasi-periodic otherwise. For a case of the former, consider $p = 0$, which gives us

$$\theta = \Big(\frac{n}{2\pi}\Big)\arctan\Big(\frac{\sin 2\pi/n}{1+\cos 2\pi/n}\Big) = \frac{1}{2};$$

hence periodic orbits.

Fig. 6. A periodic orbit on the left with the full attractor on the right.

3) Equidistribution over the attractor: Put $m = 2$ and $C = \{(1,0),(0,1),(2,3)\}$. We set $p = 1/4$. For any $v \in V$, we have

$$\lambda_v = p + \frac{1-p}{3}\big(\omega^{v_1} + \omega^{v_2} + \omega^{2v_1+3v_2}\big).$$

We found numerically $W = \{(1,0),(1,-1),(-1,0),(-1,1)\}$ and $\vartheta = \{\theta_1, \theta_2\}$, where

$$\begin{cases} \theta_1 = \big(\frac{n}{2\pi}\big)\arctan\Big(\frac{\sin 2\pi/n + \sin 4\pi/n}{2+\cos 2\pi/n + \cos 4\pi/n}\Big) \\ \theta_2 = \big(\frac{n}{2\pi}\big)\arctan\Big(\frac{-\sin 2\pi/n}{1+3\cos 2\pi/n}\Big). \end{cases}$$

By (12),

$$y_v^*(t) = \sum_{k=1,2}\Big(a_{v,\theta_k}\cos\frac{2\pi\theta_k t}{n} + b_{v,\theta_k}\sin\frac{2\pi\theta_k t}{n}\Big)x.$$

Computer experimentation points to the linear independence of the numbers $1, \theta_1, \theta_2$ over the rationals. If so, then Assumption 2.9 from Section II-D holds and, by Thm. 2.10, any limiting orbit is uniformly distributed over the attractor $\mathbb{S}$ (Fig. 7).

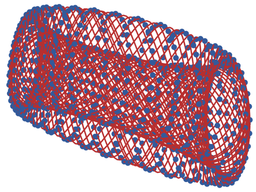

Fig. 7. A single agent's orbit is uniformly distributed around its attractor.

IV. Dynamic Social Networks

We define a *mixed* model of contrarian opinion dynamics. Let $\mathcal{M} = \{C_1, \ldots, C_s\}$ be a set of s nonempty subsets, each one spanning the vector space V. At each time step t, we define the matrix F_C by choosing, as convolution set C, a random, uniformly distributed element of $\mathcal{M}$. As before, we assume that $\mathbf{1}^\top x = 0$. Let $\lambda_{j,v}$ be the eigenvalue of F_{C_j} associated with $v \in V$. Given an infinite sequence I_∞ of indices from $[s]$, we denote by $I_t = k_1, \ldots, k_t$ be the first t indices of I_∞, and we write $\Lambda_v(I_t) = \prod_{k \in I_t} \lambda_{k,v}$. We generalize (4) into

$$x(t) = \frac{1}{N} \sum_{v \in V \setminus \{\mathbf{0}\}} \Lambda_v(I_t)\, \psi^v z_v\,, \tag{15}$$

where z_v is the row vector $\sum_{u \in V} \omega^{-\langle v,u \rangle} x_u$.

A. Spectral decomposition

Write $\lambda_v^\times = \left| \prod_{j=1}^s \lambda_{j,v} \right|^{1/s}$ and $\lambda = \max_{v \in V \setminus \{\mathbf{0}\}} \lambda_v^\times$; because all the eigenvalues other than $\lambda_{j,\mathbf{0}} = 1$ lie strictly inside the unit circle, we have $\lambda < 1$. Without loss of generality, we can assume that $\lambda > 0$. Indeed, suppose that $\lambda = 0$; then, for every $v \in V \setminus \{\mathbf{0}\}$, there is $j = j(v)$ such that $\lambda_{j,v} = 0$. This presents us with a "coupon collector's" scenario: with probability at most $N(1 - 1/s)^t \leq Ne^{-t/s}$, we have $\Lambda_v(I_t) \neq 0$ for at least one nonzero $v \in V$. In other words, with high probability, every coordinate of $x(t)$ in the eigenbasis will vanish after $O(s \log N)$ steps; hence $x(t) = 0$ for all t large enough. This case is of little interest, so we dismiss it and assume that λ is positive. We redefine $W = \{v \in V \mid \lambda_v^\times = \lambda\}$. Let $W' = \{v \in V \mid \lambda_v^\times < \lambda\}$.

Lemma 4.1. *If W' is nonempty, there exists $c < 1$ such that, with high probability, for all t large enough,*

$$\max_{w' \in W'} |\Lambda_{w'}(I_t)| \leq c^t \min_{w \in W} |\Lambda_w(I_t)|.$$

Note that the high-probability event applies to *all* times t larger than a fixed constant. The proof involves the comparison of two multiplicative random walks.

Proof. Fix $w \in W$ and $w' \in W'$. We prove that $|\Lambda_{w'}(I_t)| \leq c^t |\Lambda_w(I_t)|$. If $\lambda_{w'}^\times = 0$, then $\lambda_{j,w'} = 0$, for some j. With high probability, the sequence I_t includes the index j at least once for any t large enough; hence $|\Lambda_{w'}(I_t)| = 0$ and the lemma holds. Assume now that $\lambda_{w'}^\times > 0$; for all j, both of $\lambda_{j,w}$ and $\lambda_{j,w'}$ are nonzero. Write $S_v(I_t) = \log \prod_{k \in I_t} |\lambda_{k,v}|$, for $v = w, w'$, and note that $S_v(I_t) = t \log \lambda_v^\times + \sum_{k \in I_t} \sigma_{k,v}$, where $\sigma_{k,v} = \log|\lambda_{k,v}| - \log \lambda_v^\times$. Let $\sigma = \max_{k,v} |\sigma_{k,v}|$. The random variables $\sigma_{k,v}$ are unbiased and mutually independent in $[-\sigma, \sigma]$. Classic deviation bounds [23] give us $\mathbb{P}\Big[\big| \sum_{k \in I_t} \sigma_{k,v} \big| > b \Big] < 2e^{-b^2/(2t\sigma^2)}$. It follows that $\big| S_v(I_t) - t \log \lambda_v^\times \big| = O\big(\sigma \sqrt{t \ln(tN)}\,\big)$ with probability $1 - a/(tN)^2$, for an arbitrarily small constant $a > 0$. Since $\sum_{t>0} 1/t^2 = \pi^2/6$, it follows that, for arbitrarily small fixed $\varepsilon > 0$ and all t large enough, with probability at least $1 - \varepsilon/N^2$,

$$\begin{aligned} \log \frac{|\Lambda_w(I_t)|}{|\Lambda_{w'}(I_t)|} &= S_w(I_t) - S_{w'}(I_t) \\ &\geq t \log \frac{\lambda_w^\times}{\lambda_{w'}^\times} - O\big(\sigma \sqrt{t \log(tN)}\,\big) \geq \frac{t}{2} \log \frac{\lambda_w^\times}{\lambda_{w'}^\times}\,, \end{aligned}$$

for any given $w \in W$ and $w' \in W'$. We conclude by setting $c = \max_{w \in W, w' \in W'} \sqrt{\lambda_{w'}^\times / \lambda_w^\times}$ and using a union bound. □

We define the scaled orbit $y(t) = x(t)/\lambda^t$. Reprising the argument from Thm. 2.3, we conclude from (15) that, with high probability, the limiting orbit is of the form

$$\begin{aligned} y^*(t) &= \frac{1}{N} \sum_{h \in W} \Big(\prod_{k \in I_t} \frac{\lambda_{k,h}}{\lambda} \Big) \psi^h z_h \\ &= \frac{1}{N} \sum_{h \in W} \Big(\prod_{k \in I_t} \frac{|\lambda_{k,h}|}{\lambda} \Big)\, \omega^{\sum_{k \in I_t} \theta_{k,h}}\, \psi^h z_h, \end{aligned}$$

where $\lambda_{k,h} := |\lambda_{k,h}| \omega^{\theta_{k,h}}$. It follows that

$$y_v^*(t) = \frac{1}{N} \sum_{h \in W} \Big(\prod_{k \in I_t} \frac{|\lambda_{k,h}|}{\lambda} \Big)\, \omega^{\sum_{k \in I_t} \theta_{k,h} + \langle h,v \rangle} \sum_{u \in V} \omega^{-\langle h,u \rangle} x_u\,.$$

If we put $X_h = \frac{2\pi}{n}\big(\langle h,v \rangle + \sum_{k \in I_t} \theta_{k,h}\big)$, then, with a_h and b_h being the row vectors defined in Thm. 2.3,

$$y_v^*(t) = \frac{1}{N} \sum_{h \in W} \Big(\prod_{k \in I_t} \frac{|\lambda_{k,h}|}{\lambda} \Big) \Big((a_h x) \cos X_h + (b_h x) \sin X_h \Big).$$

B. Surprising attractors

Adding mixing to a model increases the entropy of the system. It is thus to be expected that the attractor of a mixed model should have higher dimensionality than its pure components. The surprise is that this need not be the case. We exhibit instances of contrarian opinion dynamics where mixing *decreases* the dimension of the attractor. To keep the notation simple, we consider two pure models $\mathcal{M}_1 = \{C_1\}$, $\mathcal{M}_2 = \{C_2\}$ alongside their mixture $\mathcal{M}_3 = \{C_1, C_2\}$.

Theorem 4.2. *The dimension of the mixture's attractor can be arbitrarily smaller than those of its pure components, i.e., for any $k \in [m]$, there is a choice of C_1 and C_2 such that $\dim \mathcal{M}_3 = k$ and $\dim \mathcal{M}_1 = \dim \mathcal{M}_2 = m$.*

Proof. We define $C_1 = (e_1, \ldots, e_m)$ and $C_2 = (e_1, \ldots, e_k, 2e_{k+1}, \ldots, 2e_m)$, for any $k \in [m]$, where e_i is the one-hot vector of V whose i-th coordinate is 1 and all the others 0. Let W_i be the set W corresponding to the system $\mathcal{M}_i$. We easily verify that $W_1 = \pm C_1$ and $W_2 = \pm\big\{e_1, \ldots, e_k, 2^{-1}e_{k+1}, \ldots, 2^{-1}e_m\big\}$, where 2^{-1} is the inverse of 2 in the field $\mathbb{Z}/n\mathbb{Z}$. A vector $v \in W_i$ and its negative contribute to the same ellipse, so we have $\dim \mathcal{M}_1 = \dim \mathcal{M}_2 = m$. We note that $|\lambda_{k,h}| = \lambda = \big|1 + \frac{1-p}{m}(\omega - 1)\big|$, for $h \in W_1 \cup W_2$;

hence $\lambda_v^\times = \lambda$ for $h \in W_1 \cap W_2$ and $\lambda_v^\times < \lambda$ for all other values of h. It follows that $\dim \mathcal{M}_3 = k$. □

Fig. 8 illustrates Thm. 4.2. for $m = 2$, $n = 29$, $C_1 = \{(1,0),(0,1)\}$ and $C_2 = \{(1,0),(0,2)\}$. The initial positions are random and identical in all three cases.

Fig. 8. The two attractors of the pure models $\mathcal{M}_i$ $(i = 1, 2)$ on the left, with the lower-dimensional attractor of the mixture on the right.

We can generalize the mixed model by picking C_1 (resp. C_2) with probability $1 - q$ (resp. q), where $0 \le q \le 1$. For this, we redefine $\lambda_v^\times(q) = \left|\lambda_{1,v}^{1-q}\lambda_{2,v}^{q}\right|$ and $\lambda(q) = \max_{v \in V\setminus\{\mathbf{0}\}} \lambda_v^\times(q)$.

Theorem 4.3. *The mixture's attractor can be larger than those of its pure components, i.e., there is a choice of C_1 and C_2 such that* $\dim \mathcal{M}_3 > \dim \mathcal{M}_1 = \dim \mathcal{M}_2 = m$.

Proof. Borrowing the notation of the previous proof, we define $C_1 = (e_1, \dots, e_m)$ and $C_2 = (2e_1, \dots, 2e_m)$ and verify that $W_1 = \pm C_1$ and $W_2 = \pm\{2^{-1}e_1, \dots, 2^{-1}e_m\}$; hence $\dim \mathcal{M}_1 = \dim \mathcal{M}_2 = m$. Assuming that $n > 3$, we note that the sets W_1 and W_2 are disjoint. Regarding the mixed system, we have $W(q) = \{v \in V : \lambda_v^\times(q) = \lambda(q)\}$, where $W(0) = W_1$ and $W(1) = W_2$. Around $q = 0$, we have, for all $v \in W(0)$,

$$\lambda_v^\times(q) = \left|1 + \frac{1-p}{m}(\omega - 1)\right|^{1-q} \times \left|1 + \frac{1-p}{m}(\omega^2 - 1)\right|^{q}. \quad (16)$$

Since $W(0) \neq W(1)$, by continuity, there are $q \in (0,1)$, $w \in W(q) \setminus W(0)$ such that $\lambda_w^\times(q)$ is equal to the r.h.s. of (16). This implies that $W(q) \supseteq W(0) \cup \{w\}$, which yields Thm. 4.3. □

Fig. 9 illustrates Thm. 4.3 for $m = 2$, $n = 29$, $p = 0.9$, $q = 0.0306$, $C_1 = \{(1,0),(0,1)\}$ and $C_2 = \{(2,0),(0,2)\}$. The initial positions are random and identical in all three cases.

Fig. 9. The two attractors of the pure models $\mathcal{M}_i$ $(i = 1, 2)$ on the left, with the higher-dimensional attractor of the mixture on the right.

V. CONCLUSIONS

We studied an opinion dynamics model where agents update their opinions at discrete time steps. The updates occur through a weighted average of each agent's own opinion and those of its neighbors in a social network in the form of a Cayley graph.

In this model, agents either converge to a nonconsensual equilibrium or exhibit cyclic behavior around an attractor, with the attractor's geometry being dependent on the structure of the Cayley graph. We classified different types of attractors, analyzed the agents' velocities and distribution on the attractor, and studied the effects of a randomly changing network.

References

[1] C. Bernardo, C. Altafini, A. Proskurnikov, and F. Vasca, "Bounded confidence opinion dynamics: A survey," *Automatica*, vol. 159, p. 111302, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109823004661

[2] F. Fagnani and P. Frasca, *Introduction to Averaging Dynamics over Networks*, 1st ed. Springer, 2018, vol. 472.

[3] C. F. Alford, *Group Psychology and Political Theory*. Yale University Press, 1994. [Online]. Available: http://www.jstor.org/stable/j.ctt1dt007c

[4] C. Heese, "Information Frictions and Opposed Political Interests," *American Economic Association*, 2022.

[5] S. Galam and T. Cheon, "Asymmetric Contrarians in Opinion Dynamics," *Entropy*, vol. 22, no. 1, p. 25, 2020. [Online]. Available: https://doi.org/10.3390/e22010025

[6] S. Galam, "From 2000 Bush–Gore to 2006 Italian elections: voting at fifty-fifty and the contrarian effect," *Quality & Quantity*, vol. 41, no. 4, pp. 579–589, 2007.

[7] J. R. P. French, "A Formal Theory of Social Power," *Psychological Review*, vol. 63, no. 3, pp. 181–194, 1956.

[8] M. H. DeGroot, "Reaching a Consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974. [Online]. Available: http://www.jstor.org/stable/2285509

[9] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *The Journal of Mathematical Sociology*, vol. 15, no. 3-4, pp. 193–206, 1990. [Online]. Available: https://doi.org/10.1080/0022250X.1990.9990069

[10] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Adv. Complex Syst.*, vol. 3, no. 1-4, pp. 87–98, 2000. [Online]. Available: https://doi.org/10.1142/S0219525900000078

[11] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis, "On Krause's Multi-Agent Consensus Model With State-Dependent Connectivity," *IEEE Trans. Autom. Control.*, vol. 54, no. 11, pp. 2586–2597, 2009. [Online]. Available: https://doi.org/10.1109/TAC.2009.2031211

[12] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence: models, analysis and simulation," *J. Artif. Soc. Soc. Simul.*, vol. 5, no. 3, 2002. [Online]. Available: http://jasss.soc.surrey.ac.uk/5/3/2.html

[13] A. Bhattacharyya, M. Braverman, B. Chazelle, and H. L. Nguyen, "On the convergence of the Hegselmann-Krause system," in *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, R. D. Kleinberg, Ed. ACM, 2013, pp. 61–66. [Online]. Available: https://doi.org/10.1145/2422436.2422446

[14] S. R. Etesami and T. Başar, "Game-Theoretic Analysis of the Hegselmann-Krause Model for Opinion Dynamics in Finite Dimensions," *IEEE Trans. Autom. Control.*, vol. 60, no. 7, pp. 1886–1897, 2015. [Online]. Available: https://doi.org/10.1109/TAC.2015.2394954

[15] A. Martinsson, "An improved energy argument for the Hegselmann-Krause model," *Journal of Difference Equations and Applications*, vol. 22, no. 4, pp. 513–518, 2016.

[16] K. Eekhoff, "Opinion formation dynamics with contrarians and zealots," *SIAM J. Undergraduate Research Online*, vol. 12, 2019.

[17] H. Ferraz de Arruda, A. Benatti, F. N. Silva, C. H. Comin, and L. da Fontoura Costa, "Contrarian effects and echo chamber formation in opinion dynamics," *Journal of Physics: Complexity*, vol. 2, no. 2, p. 025010, mar 2021. [Online]. Available: https://dx.doi.org/10.1088/2632-072X/abe561

[18] R. Muslim, M. J. Kholili, and A. R. Nugraha, "Opinion dynamics involving contrarian and independence behaviors based on the Sznajd model with two-two and three-one agent interactions," *Physica D: Nonlinear Phenomena*, vol. 439, p. 133379, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167278922001373

[19] B. Chazelle, K. Karntikoon, and J. Nogler, "The Geometry of Cyclical Social Trends," *arXiv preprint arXiv:2403.06376*, 2024. [Online]. Available: https://arxiv.org/abs/2403.06376

[20] E. Seneta, *Non-negative Matrices and Markov Chains*, 2nd ed., ser. Springer Series in Statistics. Springer, 2006.

[21] L. Kuipers and H. Niederreiter, *Uniform Distribution of Sequences*, ser. A Wiley-Interscience publication. Wiley, 1974.

[22] J. Cassels, *An Introduction to Diophantine Approximation*. Cambridge University Press, 1957.

[23] N. Alon and J. H. Spencer, *The Probabilistic Method, Third Edition*, ser. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 2008.