MINIMIZING POLARIZATION IN GEOMETRIC OPINION DYNAMICS ON SOCIAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose a multidimensional geometric opinion dynamics model on social networks, in which opinions evolve through alignment-based interactions. We rigorously analyze the dynamics to identify its characteristics and show that polarization can emerge endogenously. To mitigate the effect of polarization, we introduce recommendation-based interventions with exposure as a controllable parameter, yielding a nonlinear optimization problem. We show that the dynamics have a Lipschitz-continuous Jacobian and that a broad class of resulting objectives has Lipschitz-continuous gradients, establishing their tractability for optimization. We further apply a Zeroth-order projected optimization method to minimize polarization using only function evaluations, with guarantees of convergence to first-order stationary points up to smoothing bias. Experiments on synthetic and real-world network topologies demonstrate that recommendations effectively reduce polarization, and that the zeroth-order method achieves performance comparable to first-order approaches.

1 INTRODUCTION

Opinion dynamics studies how beliefs evolve through repeated interactions among individuals on social networks and exposure to external information. In these volatile and coupled systems, striking outcomes like polarization and consensus are often observed, where extreme groups are formed by individuals with extremely aligned or anti-aligned opinions. Such extreme cases give rise to echo chambers and bubbles, which have crucial socio-economic consequences (McCoy et al., 2018).

Polarization has been thoroughly studied in the form of geometric opinion dynamics, where opinions are modeled as multidimensional unit vectors with each coordinate representing a different topic. In this dynamics, opinion vectors react to other vectors based on their alignments. Namely, positive alignment pulls them closer, while negative alignment repels. In its basic form, an opinion $u^{(k)} \in \mathbb{S}^{d-1}$ interacting with an external signal $v^{(k)} \in \mathbb{S}^{d-1}$ evolves following (Hązła et al., 2019),

$$u^{(k+1)} = \frac{u'^{(k)}}{\|u'^{(k)}\|}, \qquad u'^{(k)} = u^{(k)} + \eta \, \langle u^{(k)}, v^{(k)} \rangle \, v^{(k)}. \tag{1}$$

where η is the scaling factor of influence. The push-pull mechanism in this dynamics makes polarization an endogenous consequence. This contrasts with classical averaging-based models such as DeGroot (1974) and Friedkin & Johnsen (1990), where even disagreeing neighbours pull opinions toward their weighted average, and remaining disagreement is caused by stubbornness rather than interactions. But in geometric opinion dynamics, agreement attracts while disagreement repels, allowing polarization to arise from the interaction rule itself. Existing works have strong polarization results, but they primarily focus on interaction in simplified settings, such as between pairs of opinions or under global influence. However, in a more realistic setting, opinions correspond to individuals embedded in a social network that governs their interactions. It fundamentally determines how the overlapping influences update opinion locally and aggregate across society, making the emergence of special cases like polarization nontrivial and harder to analyze. While the existing literature focuses either on social networks or the geometry of multidimensional opinion updates, in this work, we aim to unify both aspects into a single model. This raises a natural question: how does polarization emerge from the interaction between opinion geometry and network structure, and can such dynamics be systematically influenced to minimize it through external interventions?

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

To answer this, our first contribution is a model that combines geometric opinion dynamics and social networks (Section 2), where each opinion evolves through the aggregated alignment and weighted influence of its neighbors. The dependence on both alignment and topology makes the dynamics nonlinear and difficult to characterize. Multiple neighboring influences can reinforce, cancel, or redirect each other, making each update strongly dependent on the surrounding opinion configuration. To characterize this model, we rigorously analyze the dynamics under symmetric interactions and show that, despite the nonlinear and state-dependent updates, the alignment energy converges monotonically and every accumulation point of the trajectory is a fixed point (Theorem 2.1). We further illustrate the fixed-point landscape through both polarized and non-polarized configurations.

To mitigate internal polarization tendencies, our second contribution is to introduce *recommendations* as external interventions (Section 3). These represent exogenous content (e.g., social media, adverts, etc.) that can be exposed to the opinions. We extend our initial model to incorporate the interventions and propose a two-step model: recommendation interaction followed by social interaction. Keeping these channels separate preserves the underlying peer network interactions and allows external influence to propagate through subsequent social interaction. In particular, changing an opinion through recommendation exposure also changes the strengths and signs of its subsequent geometric interactions with its neighbors, making the eventual effect of an intervention inherently network-mediated. This allows us to use recommendation exposure as a control parameter to minimize polarization, but the induced optimization problem is nonconvex due to the nonlinearity of the dynamics. We show that, under reasonable metric assumptions, the resulting objective is smooth with Lipschitz-continuous gradients (Theorem 3.1). Importantly, this tractability result is not specific to squared alignment: it applies to finite combinations of sufficiently smooth downstream measures, allowing the intervention objective to encode alternative notions of disagreement, diversity, or welfare. Building on this result, we utilize recommendation exposure and construct an objective function to minimize polarization along with distribution drift. Finally, we apply a Zeroth-order method (Algorithm 1) that optimizes the objective using only function evaluations (Section 4). We also present empirical results showing that the proposed dynamics can produce polarized configurations on social networks, and that recommendation interventions can reduce our polarization measure while controlling distributional drift (Section 5). Together, these results provide a unified framework for analyzing and controlling polarization in networked geometric opinion dynamics.

1.1 Related Work

Classical models of opinion dynamics study the evolution of scalar opinions on a network using averaging dynamics (DeGroot, 1974; Friedkin & Johnsen, 1990; Hegselmann & Krause, 2002). Extensions of these models include topological analysis (Friedkin et al., 2016; Olfati-Saber et al., 2007), multidimensional opinions (Parsegov et al., 2016; De Pasquale & Valcher, 2022), stubbornness (Ghaderi & Srikant, 2014), heterogeneous susceptibility (Abebe et al., 2018), and co-evolving networks (Aghbolagh et al., 2023), but are inherently contractive. In contrast, we aim to analyze the emergence of polarization in social networks. Geometric models of opinion dynamics study opinion evolution based on alignment (Hązła et al., 2019; Lohe, 2009). These models capture both attraction and repulsion in interactions and can induce polarization naturally. Recent works establish formal guarantees for broader classes of interactions (Gaitonde et al., 2021; Alidou et al., 2024). Some variations also consider recommendations as interventions (Dean & Morgenstern, 2022; Agarwal & Brown, 2023), but largely operate on simplified universal settings, focusing on recommendation accuracy and user welfare (Yao et al., 2023a;b; 2024; Kunaver & Požrl, 2017).

Another parallel line of work studies how interventions affect opinion diversity, especially through network design and link intervention (Gaitonde et al., 2020; Musco et al., 2018; Chitra & Musco, 2020; Santos et al., 2021). But they do not explain how the consequences of these interventions may propagate through connected opinions on a social network. Recent socially-aware recommender systems use network structure to mitigate opinion clusterization under recommendation feedback (Schüepp et al., 2026), while our work instead studies multidimensional geometric dynamics where peer interactions can be both attractive and repulsive. Influence-maximization methods instead optimize information spread under diffusion (Kempe et al., 2003; Chen et al., 2010; 2016; Wang & Wang, 2023), whereas we study continuous multidimensional opinion evolution on a social network. Our work bridges these directions by combining networked geometric dynamics with controllable recommendations to study and mitigate polarization.

2 GEOMETRIC OPINION DYNAMICS ON SOCIAL NETWORKS

2.1 FORMULATION

Let $\Gamma = (V, E)$ be a weighted undirected social network that connects $N = |V|$ opinions. The edges are encoded in a row-stochastic symmetric adjacency matrix with non-negative values $W \in \mathbb{R}^{N\times N}$, which belongs to the set $\mathcal{W} : \{w_{ij} \geq 0, \sum_j w_{ij} = 1, w_{ij} = w_{ji}\}$. An edge $W_{ij} > 0$ means opinion i is influenced by opinion j, and w_{ii} is the self-interaction weight. Every node i in the network holds a unit-length vector as an opinion on the unit sphere, $o_i \in \mathbb{S}^{d-1}$ that is a multidimensional message where each coordinate corresponds to a specific topic. The collection of all opinions representing the existing individuals in the social network is the opinion matrix O, which belongs to the product Riemannian manifold $O \in \mathcal{M} = (\mathbb{S}^{d-1})^N$. The opinions evolve through the geometric opinion dynamics in Equation (1) and aggregate the influence from their neighbors, where η is replaced by the edge weight w_{ij},

$$o_i^{(k+1)} = \frac{o_i'^{(k)}}{\|o_i'^{(k)}\|}, \quad \text{where,} \quad o_i'^{(k)} = o_i^{(k)} + \sum_j^N w_{ij} \cdot \langle o_i^{(k)}, o_j^{(k)} \rangle \cdot o_j^{(k)}. \tag{2}$$

Let $O^{(k)}$ be the opinion matrix at time-step k, then the opinion matrix at time-step $k+1$, $O^{(k+1)}$ is given through the following matrix form.

$$F(O^{(k)}; W) = O^{(k+1)} = \mathcal{N}(O^{(k)} + (W \odot O^{(k)} O^{(k)\top}) O^{(k)}). \tag{3}$$

where $\odot$ is the elementwise Hadamard product, and $\mathcal{N}(\cdot)$ is a row-normalization operator that projects the unnormalized opinions back onto the unit sphere $\mathbb{S}^{d-1}$. For our analysis of the model, we assume the network connections are not affected by the opinions updates, making the adjacency matrix static $W^{(k)} = W$. So, we simplify the notation to $F(O) : \mathcal{M} \to \mathcal{M}$. Next, we analyze the proposed dynamics to derive how it evolves in terms of pairwise alignment.

2.2 MEASURING COLLECTIVE ALIGNMENT AS ENERGY

We aim to observe the evolution of the opinions when they are synchronously and continually interacting with each other. To achieve that, we measure the total alignment in the network as an energy function that calculates the aggregated sign-invariant weighted inner products between the opinions.

$$\phi(O) = \sum_{i,j}^N w_{ij} \langle o_i, o_j \rangle^2. \tag{4}$$

The energy increases with sign-invariant alignment and is bounded as $0 \leq \phi(O) \leq N$. Throughout the paper, we use the normalized energy $\phi(O)/N$ as a smooth measure of polarization, understood as directional concentration under alignment or anti-alignment. On a connected network, $\phi(O) = N$ implies that all opinions lie on a common axis, $o_i \in \{\pm v\}$, which includes consensus and bipolar configurations as extreme cases. At intermediate values, however, this measure alone does not distinguish consensus-like from bipolar structure, so we use additional geometric diagnostics when making this distinction empirically. Since normalization removes scaling, the energy updates are governed by the angular update in the tangent space. Now, the tangent space for o_i is defined as: $T_{o_i}\mathbb{S}^{d-1} = \{u \in \mathbb{R}^d : \langle u, o_i \rangle = 0\}$. We also denote the perpendicular component of a vector $x \in \mathbb{R}^d$ in terms of o_i as, $x^{\perp_i} := x - \langle x, o_i \rangle o_i$. Then, we derive the tangential or Riemannian Gradient of ϕ to determine the direction of opinion change in terms of energy.

Definition 2.1. *For a smooth function f on a sphere product manifold, the Riemannian gradient is the tangent vector that defines the direction of steepest ascent with respect to the manifold geometry (Sutti & Yueh, 2024).*

Proposition 2.1 (Riemannian Gradient)**.** *Given an energy function ϕ, the Riemannian gradient $g = \operatorname{grad} \phi(O) = (g_1, \ldots, g_N)$ on the manifold is,*

$$g_i = \operatorname{grad}_{o_i} \phi(O) = 4 \sum_{j=1}^N w_{ij} \langle o_i, o_j \rangle \, o_j^{\perp_i}, \qquad \|g\|_F^2 = \sum_i \|g_i\|^2. \tag{5}$$

Proof Sketch. For derivation, we first differentiate ϕ in terms of o_i and get the Euclidean gradient. Then we take the tangential projection of the Euclidean gradient to get the Riemannian gradient. □

With the Riemannian gradient, we can start to analyze the tangential Hessian and curvature of the energy function.

2.3 Riemannian Hessian and Curvature of Energy

For a smooth function $f(x)$ on a Riemannian manifold, the Riemannian Hessian $\mathrm{Hess}\, f(x)$ captures the local curvature of the function along tangent directions (Absil et al., 2008; Boumal, 2023). The smoothness of f is characterized by bounding the operator norm of $\mathrm{Hess}\, f(x)$ with smoothness constant L_R, which controls the deviation from the first-order approximation,

Lemma 2.1 (Riemannian smoothness bound)**.** *Let $\phi : (\mathbb{S}^{d-1})^N \to \mathbb{R}$ be a weighted alignment measurement, and let W be symmetric and row-stochastic. Then, for a self-adjoint bilinear form, the Riemannian Hessian of ϕ at O satisfies*

$$L_R\|\boldsymbol{u}\|_F^2 \geq |\,\mathrm{Hess}\,\phi(O)[\boldsymbol{u},\boldsymbol{u}]| \qquad \forall \boldsymbol{u} = [u_1, \ldots, u_N] \in T_O\mathcal{M}, \tag{6}$$

where the smoothness constant L_R admits the bound

$$L_R \leq 4 \max_{1\leq i\leq N} \left[1 - w_{ii} + \sum_{k\neq i} w_{ik} m(\langle o_i, o_k\rangle)\right] . \leq 8 - 8W_{\min}. \tag{7}$$

with $m(c) := \max(|c|, |1-2c^2|)$, and $W_{\min} := \min_i w_{ii}$ as defined in the analysis.

Proof Sketch. We bound the operator norm of the Riemannian Hessian via its quadratic form. Using the embedded submanifold formula (Absil et al., 2008; Boumal, 2023), the Hessian decomposes into three contributions: a diagonal Euclidean term, cross-interaction terms between agents, and a curvature correction arising from the spherical geometry. The diagonal term is controlled using row-stochasticity of W, the interaction terms are bounded via Cauchy-Schwarz together with the tangency conditions, and the curvature term is bounded using the same weighted alignments in ϕ. Combining these estimates yields Equation (6). Taking the supremum over $\|\mathbf{u}\|_F = 1$, the bound $m(\cdot) \leq 1$ and $\sum_{k\neq i} w_{ik} = 1 - w_{ii}$ results in the uniform bound Equation (7). □

2.4 Energy Ascent and Fixed-Point Accumulation

Next, we apply the Riemannian ascent lemma (Absil et al., 2008; Boumal, 2023) to our energy function to bound the per-step change. The key observation is that the unnormalized update part $\sum_j w_{ij}\langle o_i, o_j\rangle o_j$ geometrically contains the Riemannian gradient of the energy function as the tangential component in terms of o_i. Using this, we derive the change in energy at every update step,

Lemma 2.2 (Single-step difference in energy)**.** *Applying the Riemannian ascent lemma to $\phi(F(O))$, and $\phi(O)$ yields the inequality,*

$$\phi(F(O)) - \phi(O) \geq \sum_{i=1}^{N} \frac{\alpha_i + W_{\min}}{4(1+\alpha_i)^2}\|g_i\|^2. \tag{8}$$

Proof Sketch. We analyze the update $F(O)$ via its First-order expansion. The unnormalized update term decomposes into $\Delta_i = \alpha_i o_i + \frac{1}{4}g_i$, where the scaling amount is $\alpha_i := \sum_{j=1}^N w_{ij}\langle o_i, o_j\rangle^2$, and the tangential direction is $g_i = 4\sum_{j=1}^N w_{ij}\langle o_i, o_j\rangle o_j^{\perp_i}$. Since normalization removes the scaling component, the update effectively moves in the direction g_i. Applying the Riemannian ascent lemma then yields $\phi(F(O)) - \phi(O) \geq \sum_i \left(\eta_i - \frac{L_R}{2}\eta_i^2\right)\|g_i\|^2$, where $\eta_i = \frac{1}{4(1+\alpha_i)}$ is induced by normalization. Using $L_R \leq 8 - 8W_{\min}$ and $W_{\min} \geq 0$, we obtain a nonnegative uniform lower bound on the coefficient. □

Theorem 2.1 (Convergence and Fixed Accumulation Points)**.** *Let $W \in \mathbb{R}^{N\times N}$ be nonnegative, row-stochastic, and symmetric, and trajectory $(O_t)_{t\geq 0} \in \mathcal{M}$ be evolving under the additive normalized dynamics Equation* (3)*. Then the energy converges monotonically, and every accumulation point of the trajectory is a fixed point with property $F(\bar{O}) = \bar{O}$ and zero tangential gradient $g_i = 0; \forall i$.*

Proof Sketch. By Lemma 2.2 and the definition of α_i and g_i, if $g_i \neq 0$, then $\alpha_i > 0$, as nonzero tangential pull means there is a nonzero inner product for opinion i and some neighbour opinion. So, $\phi(O_t)$ is nondecreasing and strictly increases whenever $g(O_t) \neq 0$. Since $(\mathbb{S}^{d-1})^N$ is compact, ϕ is bounded above, $\phi(O_t)$ has to converge. We present a proof by contradiction. Let $\bar{O}$ be an accumulation point. If $g(\bar{O}) \neq 0$, then by continuity and strict ascent, there exists a neighborhood of $\bar{O}$ on which the energy increases by at least some fixed amount, contradicting convergence of $\phi(O_t)$. Hence $g(\bar{O}) = 0$, which implies $F(\bar{O}) = \bar{O}$, so every accumulation point is a fixed point, equivalently a zero-net-tangential-pull state. □

In Theorem 2.1, we can see that with continued interactions, every accumulation point of the trajectory must be a fixed point. But this does not specify which fixed points are polarized, nor the conditions for polarizing and non-polarizing initializations. Characterizing the basins of attraction that lead to polarized versus non-polarized fixed points is beyond the scope of this work. Instead, we illustrate the fixed-point landscape through two complementary examples.

A non-polarized fixed point with local energy stability on a cycle. Consider a special cycle graph C_N structure with $N \geq 5$, with nearest-neighbor interactions of opinions in $d > 2$,

$$\begin{aligned} w_{i,i+1} = w_{i,i-1} = \tfrac{1}{2}, \qquad w_{ij} = 0 \text{ otherwise}, \\ o_i = \big(\cos(i\pi/N),\, \sin(i\pi/N),\, 0, \ldots, 0\big) \in \mathbb{S}^{d-1}. \end{aligned} \tag{9}$$

This configuration is a non-polarized fixed point, as the tangential components of the neighbouring pulls cancel, yielding $g_i = 0, \forall i$. Moreover, it is a strict local maximizer of the alignment energy modulo common rotations of all opinions. We analyze this example further in supplementary material Appendix F.2.

A clustered mean-field social network. We now consider a canonical clustered mean-field social network model, in which the opinions are partitioned into communities $B_1, \ldots, B_K$, and assume that the interaction matrix has the block form

$$w_{ij} = \begin{cases} a, & i, j \in B_r \text{ for some } r, \\ b, & i \in B_r,\ j \in B_s,\ r \neq s, \end{cases} \qquad a > b > 0. \tag{10}$$

For this family of social networks, every strict local maximizer is a polarized state modulo common rotations. We present the analysis in detail in the supplementary material Appendix F.3.

These two examples highlight a thematic distinction between types of fixed points. On one hand, non-polarized fixed points exists in special structures, while on the other hand, in structured community models, the local maxima of the energy can necessarily be polarized. We use the community model example and additional empirical evidence of polarization in our simulations (presented in Appendix L.2) as motivation and present recommendation intervention methods to mitigate this effect.

3 Intervention with Recommendations

In Section 2, we showed that the energy increases monotonically and that every accumulation point of the dynamics is a fixed point. Since these fixed points may include polarized configurations, we next study whether exposure to external *recommendations* can steer the resulting dynamics. Let $R \in \mathbb{R}^{M \times d}$ be a set of M recommendations, where the j-th row $r_j^\top$ corresponds to a unit recommendation vector $r_j \in \mathbb{S}^{d-1}$. The exposure of recommendations to the opinions is represented by a separate row-stochastic connectivity matrix, $A \in \mathbb{R}^{N \times M}$, where the row-stochasticity signifies the budget for exposure of recommendations to a single opinion. Hence, it belongs to the compact and convex set $\mathcal{A} = \{a_{ij} \geq 0, \sum_j a_{ij} = 1\ \forall i\}$. An edge $a_{ij} > 0$ means opinion o_i is influenced by recommendation r_j. This leads to an extension of the original dynamics to a two-step process,

Two-step Dynamics. Introducing recommendations into the dynamics results in a two-step dynamics where the opinions evolve via recommendation interaction, followed by social interaction.

Similar to social interaction, the recommendation interaction is,

$$\hat{o}_i^{(k)} = \frac{o_i'^{(k)}}{\|o_i'^{(k)}\|}, \quad where \quad o_i'^{(k)} = o_i^{(k)} + \sum_{j=1}^{M} a_{ij} \cdot \langle o_i^{(k)}, r_j^{(k)} \rangle \cdot r_j^{(k)}. \tag{11}$$

And the matrix form,

$$G(O^{(k)}; R, A) = \hat{O}^{(k)} = \mathcal{N}(O^{(k)} + (A \odot O^{(k)} R^{(k)\top}) R^{(k)}). \tag{12}$$

Then, the opinions go through the social interactions step as described in Equations (2) and (3), where $O^{(k)}$ is replaced by $\hat{O}^{(k)}$. Let us redefine matrix function $F : \mathcal{M} \to \mathcal{M}$,

$$F(O^{(k)}; W, R, A) = \mathcal{N}(\hat{O}^{(k)} + (W \odot \hat{O}^{(k)} \hat{O}^{(k)\top}) \hat{O}^{(k)}), \quad \text{where} \quad \hat{O}^{(k)} = G(O^{(k)}; R, A). \tag{13}$$

This is the two-step dynamics that progresses the opinions, $O^{(k+1)} = F(O^{(k)}; W, R, A)$. Here, the intervention and peer interactions are kept as separate channels, which preserves the underlying peer network rather than treating recommendations as additional stubborn nodes whose weights compete with social influence.

Although the resulting dynamics is nonlinear, it still admits useful regularity and smoothness:

Lemma 3.1. *Let $O^{(k)} \in (\mathbb{S}^{d-1})^N$ be a set of opinions, $W \in \mathbb{R}^{N \times N}$ be a non-negative row-stochastic, symmetric matrix, $R^{(k)} \in (\mathbb{S}^{d-1})^M$ be a set of recommendations, and $A \in \mathbb{R}^{N \times M}$ be the connection matrix between opinions and recommendations, then the combined interaction dynamics is regular and smooth with bounds as the map $A \mapsto F(O^{(k)}; A)$ is Lipschitz continuous and has a Lipschitz-continuous Jacobian,*

$$\|F(O^{(k)}; A) - F(O^{(k)}; B)\| \leq 4\|R\|_2 \|A - B\|. \tag{14}$$

$$\|D_A F(O^{(k)}; A) - D_A F(O^{(k)}; B)\| \leq 46\|R\|_2^2 \|A - B\|. \tag{15}$$

Proof Sketch. We decompose the dynamics F as a composition of similarity, additive update, and normalization maps. We show that each component is smooth with Lipschitz-continuous derivatives. In particular, the normalization map is smooth as $\|o_i'^{(k)}\| \geq 1$, and the Frećhet derivative of the additive updates is polynomial in $(O^{(k)}, W, R, A)$. Applying standard chain rule bounds for compositions of smooth maps yields that F is Lipschitz, that depends linearly on $\|R\|_2$, and has a Lipschitz-continuous Jacobian that depends quadratically on $\|R\|_2^2$. □

As the opinion dynamics function $F(O^{(k)}; A)$ is regular and smooth, we present an objective smoothness guarantee,

Theorem 3.1 (Smoothness of Objectives)**.** *Let $F(O^{(k)}; A)$ be a regular and smooth dynamics function with Lipschitz Jacobian on a convex and compact feasible set $\mathcal{A}$ as control parameter, and $\{\varphi_1, \ldots, \varphi_P\}$ be a set of continuously differentiable metrics with Lipschitz-continuous gradients on the reachable set, where $\varphi_i : (\mathbb{S}^{d-1})^N \to \mathbb{R}$, then for fixed regularizer parameters $\{\lambda_1, \ldots, \lambda_P\}$ where $\lambda_i \in \mathbb{R}$, any objective function in the form $Q(A) = \sum_{i=1}^{P} \lambda_i \varphi_i(F(O^{(k)}; A))$ is also continuously differentiable with Lipschitz-continuous gradient on$\mathcal{A}$.*

Proof Sketch. The objective Q is a finite linear combination of smooth metrics composed with the controlled dynamics. By Lemma 3.1, the map $A \mapsto F(O^{(k)}; A)$ has a Lipschitz-continuous Jacobian on $\mathcal{A}$, while each φ_i has a Lipschitz-continuous gradient on the reachable set. The chain rule gives

$$\nabla Q(A) = \sum_{i=1}^{P} \lambda_i D_A F(O^{(k)}; A)^\top \nabla \varphi_i(F(O^{(k)}; A)). \tag{16}$$

Both factors in each summand vary Lipschitz-continuously with A on the compact feasible set, so the finite sum also has a Lipschitz-continuous gradient. □

Implications for optimization. According to Theorem 3.1, the tractability of recommendation control is not specific to any specific metric, any finite combination of sufficiently smooth downstream measures induces a smooth objective over the exposure matrix A. This implies that with a necessary projection mechanism, gradient-based First-order methods and gradient-free Zeroth-order methods apply to the objective function under standard assumptions.

Algorithm 1 Zeroth-Order Projected Descent for Recommendation Exposure

Input: $A^0 \in \mathcal{A}$, iterations T, stepsizes $\{\eta_t\}$, smoothing radius μ, directions q
for $t = 0, 1, \ldots, T-1$ **do**
  Sample $U_{t,1}, \ldots, U_{t,q} \sim \mathrm{Unif}(\mathbb{S}^{NM-1})$
  for $r = 1, \ldots, q$ **do**
    $A^+_{t,r} = A^t + \mu U_{t,r}$
    $A^-_{t,r} = A^t - \mu U_{t,r}$
    $\widehat{g}_{t,r} = \frac{NM}{2\mu}\big(J(A^+_{t,r}) - J(A^-_{t,r})\big)U_{t,r}$
  end for
  $\widehat{g}_t = \frac{1}{q}\sum_{r=1}^{q} \widehat{g}_{t,r}$
  $A^{t+1} = \Pi_{\mathcal{A}}(A^t - \eta_t \widehat{g}_t)$
end for
Output: A^T

An objective instantiation for polarization intervention To test our theoretical results through experiments, we construct a general objective using two quantities. The first is $\phi(O)$(Equation (4)), which measures polarization from how extreme the neighbour alignments are, so minimizing it promotes orthogonalization and discourages extreme alignments or anti-alignments. However, aggressive interventions can cause extreme shifts of opinions. To resolve this, we propose a drift penalty $\psi(O, O^{(0)})$ that measures how much the distribution has drifted from the baseline original distribution $O^{(0)}$,

$$\psi(O, O^{(0)}) = \frac{1}{N^2}\sum_{i,j} k(o_i, o_j) + \frac{1}{N^2}\sum_{i,j} k(o_i^{(0)}, o_j^{(0)}) - \frac{2}{N^2}\sum_{i,j} k(o_i, o_j^{(0)}). \tag{17}$$

where $k(\cdot)$ is a Gaussian (RBF) kernel, $k_\sigma(x, y) = \exp\left(-\frac{\|x-y\|_2^2}{2\sigma^2}\right)$. This drift metric is inspired by Maximum Mean Discrepancy (MMD) and adapted for opinion distribution. This encourages being faithful to the original distribution by penalizing large drifts. We combine these two metrics into a single scalar objective. We identify the exposure matrix A as our control parameter and optimize over it, while we treat the recommendation set $R^{(k)}$ as an independent environment entity. For notational simplicity, we denote $F(O^{(k)}; A)$ as the dynamics function. The objective function

$$\min_{A \in \mathcal{A}} J(A) = \phi(F(O^{(k)}; A)) + \lambda \cdot \psi(F(O^{(k)}; A), O^{(0)}). \tag{18}$$

where $\lambda \geq 0$ is a regularization parameter that controls the trade-off between alignment minimization and distributional alignment with the original.

4 Zeroth-order Optimization for Recommendation Exposure

Section 3 established that the objective Equation (18) is tractable. However, alternative settings, such as multi-step dynamics or limited observability, may make the gradients infeasible or unavailable. These practical examples motivate us to develop a Zeroth-order approach that depends solely on function evaluations of a well-structured dynamics. The zeroth-order estimator itself is standard. We use it to demonstrate that the nonlinear recommendation and peer interaction pipeline can be optimized using only objective evaluations.

Following standard randomized smoothing approaches in Zeroth-order optimization (Flaxman et al., 2004; Nesterov & Spokoiny, 2017; Ghadimi & Lan, 2013; Liu et al., 2018), we define $J_\mu(A) := \mathbb{E}_{V \sim \mathrm{Unif}(\mathbb{B}^{NM})}[J(A + \mu V)]$, with perturbation directions $U \sim \mathrm{Unif}(\mathbb{S}^{NM-1})$. The perturbed queries may lie outside of $\mathcal{A}$, but are used only for gradient estimation. The iterates themselves get projected onto the feasible set. This is well-defined in our setting because J admits a smooth extension to a neighborhood of $\mathcal{A}$. The gradient estimation,

$$\widehat{g}_{t,r} = \frac{NM}{2\mu}\big(J(A^+_{t,r}) - J(A^-_{t,r})\big)U_{t,r}. \tag{19}$$

where q is the number of perturbation samples. Under standard randomized smoothing results (Flaxman et al., 2004; Nesterov & Spokoiny, 2017; Ghadimi & Lan, 2013), this estimator satisfies

$$\mathbb{E}[\widehat{g}_t \mid A^t] = \nabla J_\mu(A^t), \quad \mathbb{E}\big[\|\widehat{g}_t - \nabla J_\mu(A^t)\|_F^2 \mid A^t\big] \le \frac{\sigma^2}{q}, \tag{20}$$

and the smoothing bias satisfies $\|\nabla J_\mu(A) - \nabla J(A)\|_F \le B_\mu; \forall A \in \mathcal{A}$, where σ^2 is the variance, and B_μ is the smoothing bias bound. These perturbations are not real-time nor applied to live opinions, but counterfactual model rollouts from the same estimated current state utilizing our proposed opinion dynamics model from Section 3.

Using this, we present Algorithm 1, a Zeroth-order projected descent method over the feasible set $\mathcal{A}$. At each step, the iterate is projected onto $\mathcal{A}$, where $\Pi_{\mathcal{A}}$ denotes projection onto the convex and compact set $\mathcal{A}$. To characterize the algorithm, we introduce the projected gradient mapping under L_J-smoothness (Ghadimi & Lan, 2013),

$$\mathcal{G}_\eta^\mu(A) = \frac{1}{\eta}\left(A - \Pi_{\mathcal{A}}(A - \eta \nabla J_\mu(A))\right). \tag{21}$$

By Theorem 3.1 and standard properties of randomized smoothing, the smoothed objective J_μ has an L_J-Lipschitz continuous gradient on $\mathcal{A}$.

Corollary 4.1 (Zeroth-order stationarity)**.** *Let $\mathcal{A} \subset \mathbb{R}^{N\times M}$ be compact and convex, let J_μ have an L_J-Lipschitz continuous gradient, and suppose the Zeroth-order estimator satisfies Equation* (20)*. Then, for $0 < \eta \le 1/(2L_J)$, standard projected Zeroth-order optimization results (Ghadimi & Lan, 2013; Liu et al., 2018) give*

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\|\mathcal{G}_\eta^\mu(A^t)\|_F^2 \le O\left(\frac{1}{\eta T}\right) + O\left(\eta \frac{\sigma^2}{q}\right). \tag{22}$$

If the smoothing bias satisfies

$$\|\nabla J_\mu(A) - \nabla J(A)\|_F \le B_\mu, \qquad \forall A \in \mathcal{A},$$

the corresponding stationarity bound for the original objective incurs an additional $O(B_\mu^2)$ term. Thus, choosing $\eta = \Theta(T^{-1/2})$ yields an $O(T^{-1/2})$ average stationarity rate for fixed q, up to the smoothing-bias term.

The parameter q controls the standard variance–evaluation trade-off: increasing q reduces estimator variance through the $1/q$ term but requires more objective evaluations. Each iteration uses $2q$ function evaluations, for a total of $2qT$ model rollouts. We examine the effects of q and the smoothing radius μ empirically in the supplementary material.

5 Experimental Results

To complement our study of opinion dynamics, we evaluate synthetic community-based networks and real-world network topologies. The synthetic networks are sampled from Stochastic Block Models (SBM) with 200–300 nodes and 10–15 communities, while real networks are sampled as connected subgraphs of 450–500 nodes from the Facebook SNAP dataset (Leskovec & Mcauley, 2012). Each node has an opinion ($d = 20$) and access to a global set of 10–15 recommendations. We compare no recommendation, a topology-based heuristic, our Zeroth-order method, and a First-order optimizer with access to gradients through the dynamics, where the objective is Equation (18).

Polarization in simulated networks As the theoretical analysis guarantees monotone energy convergence and fixed accumulation points, but does not imply polarization or consensus, we examine this behaviour empirically. We use random initiation, noisy, and community-correlated opinions on SBM and real Facebook network topologies. For every seed, the corresponding uncontrolled trajectory was run separately to convergence and reached a polarized approximate fixed point with high principal direction share and corresponding bipolar experiments also exhibit substantial disagreement. Additional diagnostics and baseline comparisons are presented in Appendix L.

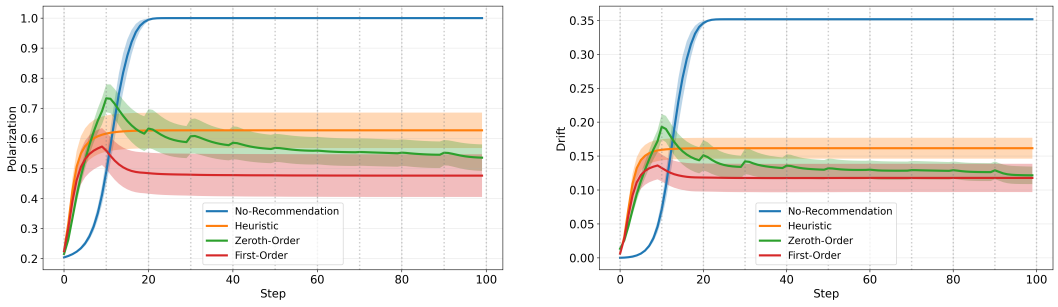

(a) Polarization (Synthetic Network)

(b) Drift (Synthetic Network)

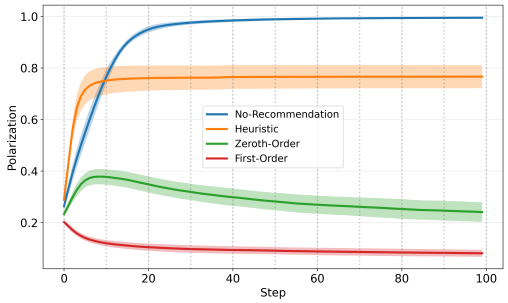

(c) Polarization (Real World Network)

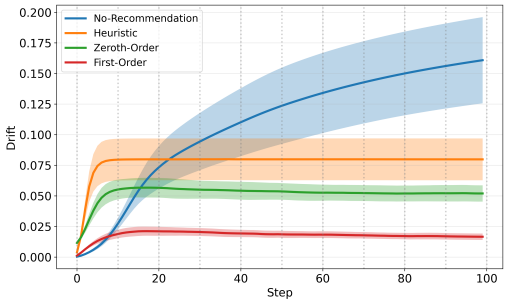

(d) Drift (Real World Network)

Figure 1: Normalized polarization, measured by the squared alignment, and drift across two settings. Solid lines denote the mean, and shaded regions the standard deviation. Polarization is highest without recommendations. Heuristic methods reduce polarization using topology, while optimization methods further improve performance, with Zeroth-order closely matching First-order.

Intervention results and algorithm comparison The results are shown in Figure 1 in population drift and normalized squared alignment over time. Without recommendations, the dynamics consistently lead to high squared alignment, consistent with the strong directional concentration discussed above. Introducing recommendations, even via heuristics, mitigates polarization by reducing squared alignment by creating a balance between intrinsic collapse and extrinsic diversification. Optimization methods further improve performance. The First-order method performs best in both tasks as it has access to the dynamics and objective gradients, while our demonstrated Zeroth-order method achieves comparable results using only function evaluations. These trends persist on real-world network topologies as well. We present additional experiments and insights in the supplementary material Appendix L that demonstrate the optimization methods in more detail.

6 Conclusion and Future Work

We introduced a networked geometric opinion-dynamics model that, unlike standard assimilative models, can capture attractive and repulsive interactions and the emergence of polarized configurations. We established monotone alignment-energy evolution and fixed-point accumulation, and showed that recommendation interventions induce a smooth optimization problem over a broad class of downstream objectives. Together, these results provide a foundation for studying how opinion geometry on networks shapes polarization and how external interventions can steer the dynamics.

Several directions follow naturally. Theoretically, an important next step is to characterize the structural conditions and basins leading to polarized versus non-polarized fixed points, and extend the analysis to asymmetric, time-varying, and heterogeneous interactions. Repeated recommendation–interaction cycles also motivate closed-loop policies that adapt exposure as opinions evolve. Empirically, validation with real opinion and recommendation data, larger networks, and alternative objectives for disagreement, polarization, diversity, and welfare would further test the framework.

AI USE STATEMENT

In this work, we used generative AI tools for writing assistance, mostly grammar and spelling. We have not used generative AI tools to generate synthetic data sets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in the writing of proofs, propose or refine hypotheses, design or provide feedback on research methodology or experiments, implement methods, assist with translation, clean and reformat datasets, support qualitative and thematic data analysis, or interpret results. We take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

We have submitted supplementary materials that include the code to generate social networks and utilize real-world topologies mentioned in the paper. The code also includes the experiments mentioned under the experiments sections in the main text and appendix. It also includes readme files to assist with reproducibility.

REFERENCES

Rediet Abebe, Jon Kleinberg, David Parkes, and Charalampos E Tsourakakis. Opinion dynamics with varying susceptibility to persuasion. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1089–1098, 2018.

P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2008.

Arpit Agarwal and William Brown. Online recommendations for agents with discounted adaptive preferences. *arXiv preprint arXiv:2302.06014*, 2023.

Hassan Dehghani Aghbolagh, Mengbin Ye, Lorenzo Zino, Zhiyong Chen, and Ming Cao. Coevolutionary dynamics of actions and opinions in social networks. *IEEE Transactions on Automatic Control*, 68(12):7708–7723, 2023.

Abdou Majeed Alidou, Júlia Baligács, Max Hahn-Klimroth, Jan Hązła, Lukas Hintze, and Olga Scheftelowitsch. Inevitability of polarization in geometric opinion exchange. *arXiv preprint arXiv:2402.08446*, 2024.

Nicolas Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.

Wei Chen, Yifei Yuan, and Li Zhang. Scalable influence maximization in social networks under the linear threshold model. In *2010 IEEE international conference on data mining*, pp. 88–97. IEEE, 2010.

Wei Chen, Tian Lin, Zihan Tan, Mingfei Zhao, and Xuren Zhou. Robust influence maximization. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 795–804, 2016.

Uthsav Chitra and Christopher Musco. Analyzing the impact of filter bubbles on social network polarization. In *Proceedings of the 13th international conference on web search and data mining*, pp. 115–123, 2020.

Giulia De Pasquale and Maria Elena Valcher. Multi-dimensional extensions of the hegselmann-krause model. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 3525–3530. IEEE, 2022.

Sarah Dean and Jamie Morgenstern. Preference dynamics under personalized recommendations. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 795–816, 2022.

Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical association*, 69(345): 118–121, 1974.

Abraham D Flaxman, Adam Tauman Kalai, and H Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. *arXiv preprint cs/0408007*, 2004.

Noah E Friedkin and Eugene C Johnsen. Social influence and opinions. *Journal of mathematical sociology*, 15(3-4):193–206, 1990.

Noah E Friedkin, Anton V Proskurnikov, Roberto Tempo, and Sergey E Parsegov. Network science on belief system dynamics under logic constraints. *Science*, 354(6310):321–326, 2016.

Jason Gaitonde, Jon Kleinberg, and Eva Tardos. Adversarial perturbations of opinion dynamics in networks. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 471–472, 2020.

Jason Gaitonde, Jon Kleinberg, and Éva Tardos. Polarization in geometric opinion dynamics. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 499–519, 2021.

Javad Ghaderi and Rayadurgam Srikant. Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate. *Automatica*, 50(12):3209–3215, 2014.

Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM journal on optimization*, 23(4):2341–2368, 2013.

Jan Hązła, Yan Jin, Elchanan Mossel, and Govind Ramnarayan. A geometric model of opinion polarization. *arXiv preprint arXiv:1910.05274*, 2019.

Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation (JASSS) vol*, 5(3), 2002.

David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146, 2003.

Matevž Kunaver and Tomaž Požrl. Diversity in recommender systems–a survey. *Knowledge-based systems*, 123:154–162, 2017.

Jure Leskovec and Julian Mcauley. Learning to discover social circles in ego networks. *Advances in neural information processing systems*, 25, 2012.

Sijia Liu, Xingguo Li, Pin-Yu Chen, Jarvis Haupt, and Lisa Amini. Zeroth-order stochastic projected gradient descent for nonconvex optimization. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pp. 1179–1183. IEEE, 2018.

M A Lohe. Non-abelian kuramoto models and synchronization. *Journal of Physics A: Mathematical and Theoretical*, 42(39):395101, 2009.

Jennifer McCoy, Tahmina Rahman, and Murat Somer. Polarization and the global crisis of democracy: Common patterns, dynamics, and pernicious consequences for democratic polities. *American behavioral scientist*, 62(1):16–42, 2018.

Cameron Musco, Christopher Musco, and Charalampos E Tsourakakis. Minimizing polarization and disagreement in social networks. In *Proceedings of the 2018 world wide web conference*, pp. 369–378, 2018.

Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17(2):527–566, 2017.

Reza Olfati-Saber, J Alex Fax, and Richard M Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

Sergey E Parsegov, Anton V Proskurnikov, Roberto Tempo, and Noah E Friedkin. Novel multidimensional models of opinion dynamics in social networks. *IEEE Transactions on Automatic Control*, 62(5):2270–2285, 2016.

Fernando P Santos, Yphtach Lelkes, and Simon A Levin. Link recommendation algorithms and dynamics of polarization in online social networks. *Proceedings of the National Academy of Sciences*, 118(50):e2102141118, 2021.

Lukas Schüepp, Carmen Amo Alonso, Florian Dörfler, and Giulia De Pasquale. Socially-aware recommender systems mitigate opinion clusterization. *arXiv preprint arXiv:2601.02412*, 2026.

Marco Sutti and Mei-Heng Yueh. Riemannian gradient descent for spherical area-preserving mappings. *arXiv preprint arXiv:2403.11726*, 2024.

Ying Wang and Yanhao Wang. Opinion-aware influence maximization in online social networks. In *2023 6th International Conference on Data Science and Information Technology (DSIT)*, pp. 214–221. IEEE, 2023.

Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. How bad is top-k recommendation under competing content creators? In *International conference on machine learning*, pp. 39674–39701. PMLR, 2023a.

Fan Yao, Chuanhao Li, Karthik Abinav Sankararaman, Yiming Liao, Yan Zhu, Qifan Wang, Hongning Wang, and Haifeng Xu. Rethinking incentives in recommender systems: are monotone rewards always beneficial? *Advances in Neural Information Processing Systems*, 36:74582–74601, 2023b.

Fan Yao, Yiming Liao, Jingzhou Liu, Shaoliang Nie, Qifan Wang, Haifeng Xu, and Hongning Wang. Unveiling user satisfaction and creator productivity trade-offs in recommendation platforms. *Advances in Neural Information Processing Systems*, 37:86958–86984, 2024.

A Appendix

This appendix provides detailed proofs and auxiliary results supporting the main text. In particular, we include the full derivations for the convergence analysis of the opinion dynamics, the regularity properties of the objective, the guarantees for the zeroth-order optimization method, and additional experiments to support our claims.

We use the same notation as in the main text. In particular, $O = (o_1, \dots, o_N) \in (\mathbb{S}^{d-1})^N$ denotes the opinion matrix, W the interaction matrix, R the recommendation matrix, and $A \in \mathcal{A}$ the exposure matrix.

B Auxiliary Calculations for Section 2

In this section, we present the geometric calculation necessary for Section 2. The key idea behind the calculations is that the dynamics in Equation (2) live on the product of spheres, so naturally the first and second-order analyses yield Riemannian gradient and Hessian rather than Euclidean ones. The Euclidean counterparts are utilized to derive Riemannian formulas using projection.

We recall the energy from Equation (4):

$$\phi(O) = \sum_{i,j=1}^{N} w_{ij} \langle o_i, o_j \rangle^2. \tag{23}$$

And define,

$$W_{\min} := \min_i w_{ii}. \tag{24}$$

For each i, the tangent space at o_i is

$$T_{o_i}\mathbb{S}^{d-1} = \{u \in \mathbb{R}^d : \langle u, o_i \rangle = 0\}, \tag{25}$$

and for any $x \in \mathbb{R}^d$ we write its component orthogonal to o_i as

$$x^{\perp i} := x - \langle x, o_i \rangle o_i. \tag{26}$$

B.1 GRADIENT FORMULA

We begin with the gradient calculation behind Proposition 2.1. The key to the derivation is that the energy couples every ordered pair (i, j), and because W is symmetric, each occurrence of o_i contributes twice to the derivative.

Lemma B.1. *Under the assumptions of Section 2, the Euclidean gradient of ϕ with respect to o_i is*

$$\nabla_{o_i}\phi(O) = 4\sum_{j=1}^{N} w_{ij}\langle o_i, o_j\rangle o_j, \tag{27}$$

and the Riemannian gradient is

$$g_i = \operatorname{grad}_{o_i}\phi(O) = 4\sum_{j=1}^{N} w_{ij}\langle o_i, o_j\rangle o_j^{\perp i}. \tag{28}$$

Proof. Fix i. Differentiating

$$\phi(O) = \sum_{j,k=1}^{N} w_{jk}\langle o_j, o_k\rangle^2. \tag{29}$$

with respect to o_i, the vector o_i appears once as o_j and once as o_k. Therefore

$$\partial_{o_i}\phi(O) = 2\sum_{j=1}^{N} w_{ij}\langle o_i, o_j\rangle o_j + 2\sum_{j=1}^{N} w_{ji}\langle o_j, o_i\rangle o_j. \tag{30}$$

Using symmetry $w_{ij} = w_{ji}$ and $\langle o_i, o_j\rangle = \langle o_j, o_i\rangle$ gives

$$\nabla_{o_i}\phi(O) = 4\sum_{j=1}^{N} w_{ij}\langle o_i, o_j\rangle o_j. \tag{31}$$

To obtain the Riemannian gradient on the sphere, we project onto the tangent space using the projection matrix $(I - o_i o_i^\top)$:

$$g_i = (I - o_i o_i^\top)\nabla_{o_i}\phi(O) = 4\sum_{j=1}^{N} w_{ij}\langle o_i, o_j\rangle o_j^{\perp i}. \tag{32}$$

This is exactly Equation (5) in the main text. □

B.2 EMBEDDED HESSIAN FORMULA

The next step is to derive the formula for calculating the Riemannian Hessian. This lets us know how much ϕ can deviate from its linear approximation along tangent directions. As the manifold in question is a product of spheres the Hessian is composed of Euclidean second-order terms and a curvature correction coming from the sphere constraint.

Lemma B.2. *Let $\bar{\phi}$ be a smooth Euclidean extension of ϕ to $(\mathbb{R}^d)^N$. For any tangent vector $\mathbf{u} = (u_1, \ldots, u_N) \in T_O\mathcal{M}$, the Riemannian Hessian,*

$$\operatorname{Hess}\phi(O)[\mathbf{u}, \mathbf{u}] = \sum_{i,j=1}^{N} u_i^\top H_{ij} u_j - \sum_{i=1}^{N}\langle o_i, \nabla_{o_i}\phi(O)\rangle\|u_i\|^2, \tag{33}$$

where $H_{ij} := \nabla^2_{o_i o_j}\bar{\phi}(O)$.

Proof. We derive the result using standard identities for Riemannian Hessians on embedded submanifolds.

Step 1: Embedded manifold identity. Let $M \subset E$ be an embedded Riemannian submanifold of a Euclidean space. A standard result (see Boumal (2023, Cor. 5.16)) states that if $\bar{G}$ is any smooth extension of $\operatorname{grad} f$, then

$$\operatorname{Hess} f(x)[u] = \operatorname{Proj}_x\big(D\bar{G}(x)[u]\big), \quad u \in T_x\mathcal{M}. \tag{34}$$

This follows from the fact that the Levi–Civita connection on M is obtained by projecting the ambient derivative onto the tangent space (see Boumal (2023, Thm. 5.9, Prop. 5.31) or Absil et al. (2008, Sec. 5.3)).

Step 2: Specialization to the sphere. For the unit sphere $\mathbb{S}^{d-1} \subset \mathbb{R}^d$, the tangent space at x is

$$T_x\mathbb{S}^{d-1} = \{v \in \mathbb{R}^d : x^\top v = 0\}, \tag{35}$$

and the orthogonal projection is

$$\mathrm{Proj}_x(v) = v - \langle x, v\rangle x. \tag{36}$$

For a function $f : \mathbb{S}^{d-1} \to \mathbb{R}$ with extension $\bar{f}$,

$$\mathrm{grad}\, f(x) = \mathrm{Proj}_x(\nabla \bar{f}(x)). \tag{37}$$

Applying the embedded Hessian formula and differentiating the projection (see Boumal (2023, Equation (7.12))) yields

$$\mathrm{Hess}\, f(x)[u] = \mathrm{Proj}_x\left(\nabla^2 \bar{f}(x)\, u\right) - \langle x, \nabla \bar{f}(x)\rangle u. \tag{38}$$

Taking the quadratic form and using $u \perp x$,

$$\mathrm{Hess}\, f(x)[u, u] = u^\top \nabla^2 \bar{f}(x)\, u - \langle x, \nabla \bar{f}(x)\rangle \|u\|^2. \tag{39}$$

Step 3: Product manifold. We now apply this identity co-ordinatewise to $\mathcal{M} = (\mathbb{S}^{d-1})^N$. Write $O = (o_1, \ldots, o_N)$ and $\mathbf{u} = (u_1, \ldots, u_N) \in T_O\mathcal{M}$. The Euclidean Hessian of $\bar{\phi}$ has block structure

$$H_{ij} := \nabla^2_{o_i o_j} \bar{\phi}(O). \tag{40}$$

Thus,

$$\mathbf{u}^\top \nabla^2 \bar{\phi}(O)\, \mathbf{u} = \sum_{i,j=1}^{N} u_i^\top H_{ij} u_j. \tag{41}$$

The curvature correction applies independently on each sphere, giving

$$\sum_{i=1}^{N} \langle o_i, \nabla_{o_i} \phi(O)\rangle \|u_i\|^2. \tag{42}$$

Combining these terms yields

$$\mathrm{Hess}\, \phi(O)[\mathbf{u}, \mathbf{u}] = \sum_{i,j=1}^{N} u_i^\top H_{ij} u_j - \sum_{i=1}^{N} \langle o_i, \nabla_{o_i} \phi(O)\rangle \|u_i\|^2. \tag{43}$$

□

The value of Lemma B.2 is conceptual: it lets us transform the Riemannian smoothness question into a matrix block calculation. Once the Euclidean Hessian blocks are known, the rest of the proof becomes estimating the quadratic forms.

Lemma B.3. *The Euclidean Hessian blocks of ϕ are given by*

$$H_{ii} = 4 \sum_{j=1}^{N} w_{ij} o_j o_j^\top, \tag{44}$$

and, for $k \neq i$,

$$H_{ik} = 4 w_{ik} \big(\langle o_i, o_k\rangle I + o_k o_i^\top\big), \tag{45}$$

where $\langle o_i, o_k\rangle = \langle o_i, o_k\rangle$. Moreover,

$$\langle o_i, \nabla_{o_i} \phi(O)\rangle = 4 \sum_{j=1}^{N} w_{ij} \langle o_i, o_j\rangle^2. \tag{46}$$

Proof. We derive the result by differentiating the Euclidean gradient of ϕ, computed in Lemma B.1:

$$\nabla_{o_i}\phi(O) = 4\sum_{j=1}^{N} w_{ij}\langle o_i, o_j\rangle\, o_j. \tag{47}$$

[Diagonal block H_{ii}] We differentiate $\nabla_{o_i}\bar{\phi}(O)$ with respect to o_i. For $j \neq i$, the term $\langle o_i, o_j\rangle o_j$ depends on o_i, while o_j is fixed. Hence

$$\nabla_{o_i}(\langle o_i, o_j\rangle o_j) = o_j o_j^\top. \tag{48}$$

The self-interaction term $j = i$ corresponds to $w_{ii}\langle o_i, o_i\rangle^2 = w_{ii}$ on the sphere, and is treated as constant in the sphere-adapted extension. Therefore, it contributes no Hessian term. Summing over $j \neq i$ yields

$$H_{ii} = \nabla^2_{o_i o_i}\bar{\phi}(O) = 4\sum_{j\neq i} w_{ij} o_j o_j^\top. \tag{49}$$

[Off-diagonal block H_{ik}, $k \neq i$] Only the term with $j = k$ depends on o_k, so we differentiate $(4w_{ik}\langle o_i, o_k\rangle o_k)$, using the product rule,

$$\nabla_{o_k}(\langle o_i, o_k\rangle o_k) = (\nabla_{o_k}\langle o_i, o_k\rangle)o_k^\top + \langle o_i, o_k\rangle I. \tag{50}$$

Since $\nabla_{o_k}\langle o_i, o_k\rangle = o_i$, this gives

$$\nabla_{o_k}(\langle o_i, o_k\rangle o_k) = o_k o_i^\top + \langle o_i, o_k\rangle I. \tag{51}$$

Thus,

$$H_{ik} = 4w_{ik}(\langle o_i, o_k\rangle I + o_k o_i^\top). \tag{52}$$

[Radial component] Taking the inner product with o_i,

$$\langle o_i, \nabla_{o_i}\phi(O)\rangle = 4\sum_{j\neq i} w_{ij}\langle o_i, o_j\rangle\langle o_i, o_j\rangle = 4\sum_{j\neq i} w_{ij}\langle o_i, o_j\rangle^2. \tag{53}$$

This completes the proof. □

C Proof of the Riemannian Smoothness Bound

We now prove the claim from the main-text lemma on Riemannian smoothness. The strategy is to write the Hessian quadratic form as three pieces and then bound each piece in a way that isolates the dependence on the local geometry $\langle o_i, o_k\rangle$.

Proof of Lemma 2.1. Let $\mathbf{u} = (u_1, \ldots, u_N) \in T_O\mathcal{M}$. Substituting the block formulas from Lemma B.3 into Lemma B.2, we evaluate the quadratic form

$$\operatorname{Hess}\phi(O)[\mathbf{u}, \mathbf{u}] = \sum_i u_i^\top H_{ii} u_i \;+\; \sum_{i\neq k} u_i^\top H_{ik} u_k \;-\; \sum_i \langle o_i, \nabla_{o_i}\phi(O)\rangle \|u_i\|^2. \tag{54}$$

Expanding each term using the expressions for H_{ii}, H_{ik}, and the radial correction yields the decomposition

$$\operatorname{Hess}\phi(O)[\mathbf{u}, \mathbf{u}] = T_1 + T_2 + T_3, \tag{55}$$

where

$$T_1 = 4\sum_{i=1}^{N}\sum_{j\neq i} w_{ij}(u_i \cdot o_j)^2, \tag{56}$$

$$T_2 = 8\sum_{i<k} w_{ik}\Big(\langle o_i, o_k\rangle(u_i \cdot u_k) + (u_i \cdot o_k)(o_i \cdot u_k)\Big), \tag{57}$$

$$T_3 = -4\sum_{i,j=1}^{N} w_{ij}\langle o_i, o_j\rangle^2 \|u_i\|^2. \tag{58}$$

The term T_1 collects the diagonal contributions $u_i^\top H_{ii} u_i$, T_2 aggregates the off-diagonal couplings $u_i^\top H_{ik} u_k$, and T_3 arises from the radial correction term in Lemma B.2.

We first bound T_1. Since $u_i \perp o_i$, for $j \neq i$ we have

$$|u_i \cdot o_j| = |u_i \cdot (o_j - \langle o_i, o_j \rangle o_i)| \leq \|u_i\| \sqrt{1 - \langle o_i, o_j \rangle^2}. \tag{59}$$

Hence

$$|T_1| \leq 4 \sum_{i=1}^{N} \left(\sum_{j \neq i} w_{ij} (1 - \langle o_i, o_j \rangle^2) \right) \|u_i\|^2. \tag{60}$$

This estimate shows explicitly why tangency matters: only the component of o_j orthogonal to o_i can interact with u_i.

Next, define

$$E_{ik} := \langle o_i, o_k \rangle (u_i \cdot u_k) + (u_i \cdot o_k)(o_i \cdot u_k). \tag{61}$$

The bound follows by restricting to the plane spanned by o_i and o_k, where the interaction depends only on $c = \langle o_i, o_k \rangle$. In this two-dimensional setting, the tangent directions admit only a few external configurations: one aligned with o_i, o_k, contributing $|c|$, and another maximizing cross-interactions, contributing $|1 - 2c^2|$. Taking the worst case gives

$$m(c) = \max\{|c|,\ |1 - 2c^2|\}, \tag{62}$$

and hence

$$|E_{ik}| \leq m(\langle o_i, o_k \rangle)\, \|u_i\| \|u_k\|. \tag{63}$$

where $m(c) := \max\{|c|, |1 - 2c^2|\} \leq 1$. Therefore

$$|T_2| \leq 8 \sum_{i<k} w_{ik} m(\langle o_i, o_k \rangle) \|u_i\| \|u_k\|. \tag{64}$$

Applying $2ab \leq a^2 + b^2$ yields

$$|T_2| \leq 4 \sum_{i=1}^{N} \left(\sum_{k \neq i} w_{ik} m(\langle o_i, o_k \rangle) \right) \|u_i\|^2. \tag{65}$$

This is the key step where the state-dependent factor $m(\langle o_i, o_k \rangle)$ appears. It reflects the fact that the cross-agent coupling is weakest in very aligned or very orthogonal configurations and strongest in intermediate geometries.

Finally,

$$|T_3| = 4 \sum_{i=1}^{N} \left(\sum_{j \neq i}^{N} w_{ij} \langle o_i, o_j \rangle^2 \right) \|u_i\|^2. \tag{66}$$

Combining the three estimates gives

$$|\operatorname{Hess} \phi(O)[\mathbf{u}, \mathbf{u}]| \leq 4 \sum_{i=1}^{N} \left[\sum_{j \neq i} w_{ij} (1 - \langle o_i, o_j \rangle^2) + \sum_{k \neq i} w_{ik} m(\langle o_i, o_k \rangle) + \sum_{j \neq i}^{N} w_{ij} \langle o_i, o_j \rangle^2 \right] \|u_i\|^2. \tag{67}$$

Since

$$\sum_{j \neq i} w_{ij} (1 - \langle o_i, o_j \rangle^2) + \sum_{j \neq i}^{N} w_{ij} \langle o_i, o_j \rangle^2 = \sum_{j \neq i}^{N} w_{ij} = 1 - w_{ii}, \tag{68}$$

we arrive at

$$|\operatorname{Hess} \phi(O)[\mathbf{u}, \mathbf{u}]| \leq 4 \sum_{i=1}^{N} \left[1 - w_{ii} + \sum_{k \neq i} w_{ik} m(\langle o_i, o_k \rangle) \right] \|u_i\|^2. \tag{69}$$

Taking the supremum over $\|\mathbf{u}\|_F = 1$ gives

$$L_R \leq 4 \max_{1 \leq i \leq N} \left[1 - w_{ii} + \sum_{k \neq i} w_{ik} m(\langle o_i, o_k \rangle) \right]. \tag{70}$$

Finally, since $m(c) \leq 1$ and $\sum_{k \neq i} w_{ik} = 1 - w_{ii}$, we obtain the uniform bound

$$L_R \leq 4 \max_i (2 - 2w_{ii}) \leq 8 - 8W_{\min}. \tag{71}$$

This proves the main-text lemma. □

D Proof of Single Step Energy Change Lemma

The proof of Lemma 2.2 is to determine the characteristic of the dynamics. In the unnormalized update Equation (2), the influence can be decomposed into a radial term, which only changes the magnitude of a row, and a tangential term, which changes the direction. Because the final normalization projects back onto the sphere, only the tangential part contributes to the First-order movement on the manifold.

Proof of Lemma 2.2. Let

$$\Delta_i := \sum_{j=1}^{N} w_{ij} \langle o_i, o_j \rangle o_j, \tag{72}$$

and decompose each neighbor o_j into its radial and tangential components relative to o_i:

$$o_j = \langle o_i, o_j \rangle o_i + o_j^{\perp_i}, \qquad o_j^{\perp_i} := o_j - \langle o_i, o_j \rangle o_i. \tag{73}$$

Substituting this into Δ_i gives

$$\Delta_i = \sum_{j=1}^{N} w_{ij} \langle o_i, o_j \rangle \left(\langle o_i, o_j \rangle o_i + o_j^{\perp_i} \right) \tag{74}$$

$$= \left(\sum_{j=1}^{N} w_{ij} \langle o_i, o_j \rangle^2 \right) o_i + \sum_{j=1}^{N} w_{ij} \langle o_i, o_j \rangle o_j^{\perp_i}. \tag{75}$$

Define

$$\alpha_i := \sum_{j=1}^{N} w_{ij} \langle o_i, o_j \rangle^2. \tag{76}$$

By Lemma B.1,

$$g_i = 4 \sum_{j=1}^{N} w_{ij} \langle o_i, o_j \rangle o_j^{\perp_i}. \tag{77}$$

Therefore,

$$\Delta_i = \alpha_i o_i + \frac{1}{4} g_i. \tag{78}$$

Therefore, the pre-normalized update is

$$o_i' = o_i + \Delta_i = (1 + \alpha_i) o_i + \frac{1}{4} g_i. \tag{79}$$

Factor out $(1 + \alpha_i)$ and define $\eta_i := \frac{1}{4(1+\alpha_i)}$. Then

$$o_i' = (1 + \alpha_i)\big(o_i + \eta_i g_i\big). \tag{80}$$

After row normalization, the scalar factor $(1 + \alpha_i)$ disappears, so the normalized update is exactly a retraction-type step in the direction g_i:

$$F_i(O) = \frac{o_i + \eta_i g_i}{\|o_i + \eta_i g_i\|}. \tag{81}$$

This decomposition is the geometric heart of the proof: the social interaction looks complicated in ambient coordinates, but on the sphere, it is simply a gradient-like step with a state-dependent step size.

Because $|\langle o_i, o_j\rangle| \le 1$ and W is row-stochastic,

$$0 \le \alpha_i \le 1, \tag{82}$$

which implies

$$\frac{1}{8} \le \eta_i \le \frac{1}{4}. \tag{83}$$

Ascent inequality. We derive the one-step increase using the Riemannian ascent lemma on the product manifold $\mathcal{M} = (\mathbb{S}^{d-1})^N$ (Absil et al., 2008). For any $O \in \mathcal{M}$ and tangent step $\Xi \in T_O\mathcal{M}$,

$$\phi(R_O(\Xi)) \ge \phi(O) + \langle \operatorname{grad}\phi(O), \Xi\rangle - \frac{L_R}{2}\|\Xi\|_F^2, \tag{84}$$

where R_O is the row-wise normalization retraction. The normalized update can be written as a retraction step

$$F(O) = R_O(\Xi(O)), \quad \Xi(O) = (\eta_1 g_1, \dots, \eta_N g_N), \quad \eta_i = \frac{1}{4(1+\alpha_i)}. \tag{85}$$

Applying the ascent lemma with $\Xi = \Xi(O)$,

$$\phi(F(O)) - \phi(O) \ge \sum_{i=1}^N \langle g_i, \eta_i g_i\rangle - \frac{L_R}{2}\sum_{i=1}^N \eta_i^2\|g_i\|^2 \tag{86}$$

$$= \sum_{i=1}^N \left(\eta_i - \frac{L_R}{2}\eta_i^2\right)\|g_i\|^2. \tag{87}$$

Now, using $L_R \le 8 - 8W_{\min}$ and representing η_i with α_i, we get,

$$\eta_i - \frac{L_R}{2}\eta_i^2 \ge \frac{1}{4(1+\alpha_i)} - \frac{8-8W_{\min}}{2}\left(\frac{1}{4(1+\alpha_i)}\right)^2 = \frac{\alpha_i + W_{\min}}{4(1+\alpha_i)^2}. \tag{88}$$

Hence

$$\phi(F(O)) - \phi(O) \ge \sum_{i=1}^N \frac{\alpha_i + W_{\min}}{4(1+\alpha_i)^2}\|g_i\|^2 \tag{89}$$

This proves Lemma 2.2.

□

E Proof of Dynamics Convergence Theorem

Theorem 2.1 is an asymptotic consequence of Lemma 2.2. Once the energy is shown to increase whenever the tangential pull is nonzero, the only possible accumulation points are those where no tangential motion remains.

Proof of Theorem 2.1. By Lemma 2.2,

$$\phi(F(O)) - \phi(O) \ge \sum_{i=1}^N \frac{\alpha_i + W_{\min}}{4(1+\alpha_i)^2}\|g_i\|^2 \ge 0. \tag{90}$$

The final inequality follows because $\alpha_i \ge 0$, $W_{\min} \ge 0$, and $\|g_i\|^2 \ge 0$ for every i. Hence the scalar sequence $(\phi(O_t))_{t\ge 0}$ is nondecreasing.

Moreover, the increase is strict whenever $g(O) \neq 0$. Indeed, if $g_i \neq 0$ for some node i, then necessarily $\alpha_i > 0$. To see this, recall that

$$\alpha_i = \sum_{j=1}^{N} w_{ij} \langle o_i, o_j \rangle^2, \tag{91}$$

and

$$g_i = 4 \sum_{j=1}^{N} w_{ij} \langle o_i, o_j \rangle o_j^{\perp i}. \tag{92}$$

Since every term in the definition of α_i is nonnegative, $\alpha_i = 0$ implies that $w_{ij}\langle o_i, o_j\rangle^2 = 0$ for every j. Consequently, $w_{ij}\langle o_i, o_j\rangle = 0$ for every j, which implies $g_i = 0$. Therefore, by contraposition,

$$g_i \neq 0 \quad \Longrightarrow \quad \alpha_i > 0. \tag{93}$$

Thus, whenever $g_i \neq 0$, the corresponding term

$$\frac{\alpha_i + W_{\min}}{4(1+\alpha_i)^2} \|g_i\|^2 \tag{94}$$

is strictly positive, even when $W_{\min} = 0$. Hence

$$g(O) \neq 0 \quad \Longrightarrow \quad \phi(F(O)) > \phi(O). \tag{95}$$

Since $\mathcal{M} = (\mathbb{S}^{d-1})^N$ is compact and ϕ is continuous, ϕ is bounded above, so $\phi(O_t)$ converges.

Let $\bar{O}$ be an accumulation point of (O_t). Suppose, for contradiction, that $g(\bar{O}) \neq 0$. Then there exists some node $i^\star$ such that $g_{i^\star}(\bar{O}) \neq 0$. As shown above, this also implies $\alpha_{i^\star}(\bar{O}) > 0$.

By continuity of $g_{i^\star}$ and $\alpha_{i^\star}$, there exist a neighborhood U of $\bar{O}$ and constants $\delta > 0$ and $a > 0$ such that

$$\|g_{i^\star}(O)\|^2 \geq \delta, \qquad \alpha_{i^\star}(O) \geq a, \qquad \forall O \in U. \tag{96}$$

Since $0 \leq \alpha_i \leq 1$, we also have $(1+\alpha_i)^2 \leq 4$. Therefore, for every $O \in U$,

$$\phi(F(O)) - \phi(O) \geq \frac{\alpha_{i^\star} + W_{\min}}{4(1+\alpha_{i^\star})^2} \|g_{i^\star}\|^2 \tag{97}$$

$$\geq \frac{\alpha_{i^\star}}{16} \|g_{i^\star}\|^2 \tag{98}$$

$$\geq \frac{a\delta}{16} > 0. \tag{99}$$

Thus every visit to U forces the energy to increase by at least the fixed positive amount $a\delta/16$. Since $\bar{O}$ is an accumulation point, infinitely many iterates enter U, which contradicts convergence of $\phi(O_t)$. Therefore $g(\bar{O}) = 0$.

We now show that $g(\bar{O}) = 0$ implies $F(\bar{O}) = \bar{O}$. From the decomposition in the proof of Lemma 2.2,

$$\Delta_i = \alpha_i o_i + \frac{1}{4} g_i. \tag{100}$$

If $g_i(\bar{O}) = 0$, then $\Delta_i = \alpha_i \bar{o}_i$, so the pre-normalized update satisfies

$$\bar{o}_i' = (1+\alpha_i)\bar{o}_i. \tag{101}$$

Normalizing this vector returns $\bar{o}_i$ itself. Hence $F(\bar{O}) = \bar{O}$, and every accumulation point is a fixed point. So, the dynamics can only accumulate at states when the net tangential influence at every node is zero. □

F Additional examples: non-polarized and polarized fixed points

F.1 Second-order variation under normalized perturbations

We first record the second-order expansion of ϕ under a normalized perturbation at a single node.

Let

$$\phi(O) = \sum_{i,j} W_{ij} \langle o_i, o_j \rangle^2, \qquad O = (o_1, \dots, o_N) \in (\mathbb{S}^{d-1})^N. \tag{102}$$

Consider perturbing node i along a tangent direction $h_i \perp o_i$:

$$o_i^\epsilon = \frac{o_i + \epsilon h_i}{\|o_i + \epsilon h_i\|}. \tag{103}$$

We derive the expansion by noting that only terms involving node i change. Using symmetry of W,

$$\phi(O^\epsilon) - \phi(O) = 2 \sum_{j=1}^{N} W_{ij} \Big(\langle o_i^\epsilon, o_j \rangle^2 - \langle o_i, o_j \rangle^2 \Big). \tag{104}$$

Using the normalization expansion

$$o_i^\epsilon = o_i + \epsilon h_i - \tfrac{1}{2}\epsilon^2 \|h_i\|^2 o_i + O(\epsilon^3), \tag{105}$$

we obtain

$$\langle o_i^\epsilon, o_j \rangle = \langle o_i, o_j \rangle + \epsilon \langle h_i, o_j \rangle - \tfrac{1}{2}\epsilon^2 \|h_i\|^2 \langle o_i, o_j \rangle + O(\epsilon^3). \tag{106}$$

Expanding the square and keeping terms up to second order,

$$\langle o_i^\epsilon, o_j \rangle^2 = \langle o_i, o_j \rangle^2 + 2\epsilon \langle o_i, o_j \rangle \langle h_i, o_j \rangle + \epsilon^2 \Big(\langle h_i, o_j \rangle^2 - \|h_i\|^2 \langle o_i, o_j \rangle^2 \Big) + O(\epsilon^3). \tag{107}$$

Substituting and grouping terms gives

$$\phi(O^\epsilon) - \phi(O) = 4\epsilon \sum_j W_{ij} \langle o_i, o_j \rangle \langle h_i, o_j \rangle + 2\epsilon^2 \sum_j W_{ij} \Big(\langle h_i, o_j \rangle^2 - \|h_i\|^2 \langle o_i, o_j \rangle^2 \Big) + O(\epsilon^3). \tag{108}$$

Defining

$$M_i := \sum_{j=1}^{N} W_{ij} o_j o_j^\top, \tag{109}$$

we obtain

$$\phi(O^\epsilon) - \phi(O) = 4\epsilon \, h_i^\top M_i o_i + 2\epsilon^2 \Big(h_i^\top M_i h_i - \langle o_i, M_i o_i \rangle \|h_i\|^2 \Big) + O(\epsilon^3). \tag{110}$$

At a local maximizer, the first-order variation must vanish along all tangent directions $h_i \perp o_i$. Since the linear term is $4\, h_i^\top M_i o_i$, this implies $M_i o_i$ is colinear with o_i, i.e., o_i is an eigenvector of M_i. Therefore,

$$\phi(O^\epsilon) - \phi(O) = 2\epsilon^2 \Big(h_i^\top M_i h_i - \langle o_i, M_i o_i \rangle \|h_i\|^2 \Big) + O(\epsilon^3). \tag{111}$$

In particular, at a strict local maximizer of ϕ, o_i must be the unique top eigenvector of M_i for every i. If the top eigenspace had dimension greater than one, one could choose a direction $h_i \perp o_i$ within that eigenspace, yielding zero second variation and contradicting strict maximality.

F.2 A NON-POLARIZED CYCLE FIXED POINT WITH LOCAL ENERGY STABILITY

We now study the non-polarized cycle configuration introduced in Section 2.4 in more detail. In addition to showing that it is a fixed point, we show that it is a strict local maximizer of the alignment energy modulo the natural invariance to common rotations of all opinions.

Consider the cycle graph C_N, with $N \geq 5$, and weights

$$w_{i,i+1} = w_{i,i-1} = \frac{1}{2}, \qquad w_{ij} = 0 \text{ otherwise}, \tag{112}$$

where indices are interpreted modulo N. Let

$$q := \frac{\pi}{N}, \qquad o_i = \big(\cos(iq), \sin(iq), 0, \dots, 0 \big) \in \mathbb{S}^{d-1}. \tag{113}$$

The opinions therefore lie in a common two-dimensional plane $P \subset \mathbb{R}^d$ and rotate by an angle q between consecutive nodes. Since $q \in (0, \pi/2)$, this configuration is not polarized.

Fixed point. Let

$$c := \cos q. \tag{114}$$

For every non-wraparound edge, $\langle o_i, o_{i+1}\rangle = c$, while across the wraparound edge the inner product is $-c$. In both cases, the sign of the inner product is compensated by the corresponding orientation of the neighboring opinion. Consequently,

$$\sum_{j=1}^{N} w_{ij}\langle o_i, o_j\rangle o_j = c^2 o_i, \qquad \forall i. \tag{115}$$

Thus the pre-normalized update is

$$o_i' = (1 + c^2)o_i, \tag{116}$$

and row normalization returns o_i. Hence

$$F(O) = O. \tag{117}$$

Equivalently, $g_i = 0$ for every node.

General tangent perturbations. We next consider an arbitrary tangent perturbation $h = (h_1, \ldots, h_N) \in T_O\mathcal{M}$. Let

$$t_i := \big(-\sin(iq), \cos(iq), 0, \ldots, 0\big) \tag{118}$$

denote the unit tangent direction within the cycle plane. Every tangent vector $h_i \perp o_i$ admits the orthogonal decomposition

$$h_i = a_i t_i + z_i, \qquad z_i \in P^{\perp}. \tag{119}$$

We perturb each opinion by row normalization,

$$o_i^{\epsilon} := \frac{o_i + \epsilon h_i}{\|o_i + \epsilon h_i\|}. \tag{120}$$

The in-plane and out-of-plane components decouple to second order because z_i is orthogonal to every vector in P. We analyze the two parts separately.

In-plane perturbations. First suppose $z_i = 0$. To second order, the perturbation is equivalently an angular perturbation

$$o_i^{\epsilon} = \big(\cos(iq + \epsilon a_i), \sin(iq + \epsilon a_i), 0, \ldots, 0\big) + O(\epsilon^2) \tag{121}$$

with the same second variation. Since each undirected cycle edge contributes once to the energy,

$$\phi(O^{\epsilon}) = \sum_{i=1}^{N} \cos^2\left(q + \epsilon(a_{i+1} - a_i)\right), \qquad a_{N+1} := a_1, \tag{122}$$

where the identity also holds for the wraparound edge because $\cos^2(\theta)$ is π-periodic.

Using

$$\frac{d}{d\theta}\cos^2\theta = -\sin(2\theta), \qquad \frac{d^2}{d\theta^2}\cos^2\theta = -2\cos(2\theta), \tag{123}$$

and

$$\sum_{i=1}^{N}(a_{i+1} - a_i) = 0, \tag{124}$$

the first-order terms cancel, and we obtain

$$\phi(O^{\epsilon}) - \phi(O) = -\epsilon^2\cos(2q)\sum_{i=1}^{N}(a_{i+1} - a_i)^2 + O(\epsilon^3). \tag{125}$$

Since $N \geq 5$,

$$\cos(2q) = \cos\left(\frac{2\pi}{N}\right) > 0. \tag{126}$$

Therefore the second variation is strictly negative unless

$$a_1 = \cdots = a_N. \tag{127}$$

The remaining zero direction corresponds exactly to rotating all opinions by the same angle, which leaves every pairwise inner product unchanged.

Out-of-plane perturbations. Now suppose $a_i = 0$, so $h_i = z_i \in P^\perp$. Define

$$s_i = \begin{cases} 1, & i = 1, \dots, N-1, \\ -1, & i = N, \end{cases} \tag{128}$$

so that

$$\langle o_i, o_{i+1} \rangle = s_i c, \qquad o_{N+1} := o_1. \tag{129}$$

Expanding the normalized perturbation gives

$$o_i^\epsilon = o_i + \epsilon z_i - \frac{\epsilon^2}{2} \|z_i\|^2 o_i + O(\epsilon^3). \tag{130}$$

Because $z_i \perp P$,

$$\langle z_i, o_j \rangle = 0 \qquad \forall i, j. \tag{131}$$

Hence the second-order change in energy is

$$\phi(O^\epsilon) - \phi(O) = -2\epsilon^2 \sum_{i=1}^{N} \left(c^2 \|z_i\|^2 - c\, s_i \langle z_i, z_{i+1} \rangle \right) + O(\epsilon^3). \tag{132}$$

To determine the sign of this quadratic form, choose any orthonormal basis of $P^\perp$ and decompose z_i coordinatewise. For each scalar coordinate b_i, the corresponding quadratic form is

$$b^\top L_s b = \sum_{i=1}^{N} \left(c^2 b_i^2 - c\, s_i b_i b_{i+1} \right). \tag{133}$$

Because

$$\prod_{i=1}^{N} s_i = -1, \tag{134}$$

a diagonal sign transformation converts the signed cycle into a cycle with the anti-periodic boundary condition

$$b_{N+1} = -b_1. \tag{135}$$

Its Fourier frequencies are

$$\theta_k = \frac{(2k+1)\pi}{N}, \qquad k = 0, \dots, N-1, \tag{136}$$

and the corresponding eigenvalues are

$$\lambda_k = c^2 - c \cos \theta_k. \tag{137}$$

Since $c = \cos(\pi/N)$,

$$\lambda_k \geq 0, \tag{138}$$

with equality only for

$$\theta_k = \frac{\pi}{N} \quad \text{or} \quad 2\pi - \frac{\pi}{N}. \tag{139}$$

Thus the real kernel is two-dimensional and is spanned by

$$b_i = \cos(iq), \qquad b_i = \sin(iq). \tag{140}$$

Consequently, the out-of-plane second variation is nonpositive and is strictly negative away from these two modes.

These zero modes also have a geometric interpretation. For any fixed $u \in P^\perp$, perturbations of the form

$$z_i = \big(\beta_1 \cos(iq) + \beta_2 \sin(iq) \big) u \tag{141}$$

are precisely the infinitesimal perturbations obtained by applying a common rotation that tilts the plane P toward u. Therefore they do not change any pairwise inner products to second order.

Local maximality modulo global rotations. Combining Equation (125) and Equation (132), for an arbitrary tangent perturbation $h_i = a_i t_i + z_i$,

$$\phi(O^\epsilon) - \phi(O) = -\epsilon^2 \cos(2q) \sum_{i=1}^{N} (a_{i+1} - a_i)^2 - 2\epsilon^2 \sum_{i=1}^{N} \left(c^2 \|z_i\|^2 - c\, s_i \langle z_i, z_{i+1} \rangle \right) + O(\epsilon^3). \tag{142}$$

Both quadratic terms are nonpositive in the energy variation. Equality occurs only along infinitesimal common rotations of the entire configuration.

Indeed, the kernel contains one in-plane rotational direction and two plane-tilting directions for each of the $d-2$ orthogonal dimensions, for a total of

$$1 + 2(d-2) = 2d - 3 \tag{143}$$

zero directions. This is exactly the dimension of the orbit of the two-dimensional configuration under common orthogonal rotations,

$$\dim O(d) - \dim O(d-2) = 2d - 3. \tag{144}$$

Thus there are no additional flat directions: the Hessian of ϕ is negative definite in every direction transverse to the global-rotation orbit.

Therefore, although the cycle configuration is not polarized, it is a strict local maximizer of the alignment energy modulo common rotations of all opinions. In this sense, the example demonstrates that non-polarized fixed points can possess genuine local energy stability rather than arising only as fragile cancellation configurations.

F.3 Clustered mean-field interaction model

We analyze the clustered mean-field model introduced in Equation (10). Let the nodes be partitioned into communities $B_1, \ldots, B_K$, and assume

$$w_{ij} = \begin{cases} a, & i, j \in B_r \text{ for some } r, \\ b, & i \in B_r,\ j \in B_s,\ r \neq s, \end{cases} \qquad a > b > 0. \tag{145}$$

Define

$$X = \sum_{i=1}^{N} o_i o_i^\top, \qquad X_r = \sum_{i \in B_r} o_i o_i^\top. \tag{146}$$

For any node $i \in B_r$, the local covariance matrix is

$$M_i = \sum_{j=1}^{N} w_{ij}\, o_j o_j^\top = aX_r + b(X - X_r) = bX + (a-b)X_r. \tag{147}$$

Thus M_i is identical for all nodes in the same block B_r.

Block collapse. At a strict local maximizer modulo common rotations of ϕ, the second variation condition implies that for each i, o_i is the unique top eigenvector of M_i. For $i \in B_r$, the matrix M_i is identical, so all nodes in B_r share the same matrix $M_r := M_i$. Let v_r denote a unit vector spanning the top eigenspace of M_r. Since o_i is a top eigenvector of M_r for every $i \in B_r$, and the top eigenspace is one-dimensional,

$$o_i \in \{\pm v_r\}, \qquad i \in B_r. \tag{148}$$

Global alignment. Write $o_i = s_i v_r$ for $i \in B_r$, where $s_i \in \{\pm 1\}$. Consider the coherent signed perturbation

$$o_i^\epsilon = s_i \frac{v_r + \epsilon h}{\|v_r + \epsilon h\|}, \qquad h \perp v_r, \quad i \in B_r. \tag{149}$$

For this coherent rotation of block B_r,

$$X_r^\epsilon = n_r v_\epsilon v_\epsilon^\top, \qquad v_\epsilon = \frac{v_r + \epsilon h}{\|v_r + \epsilon h\|}, \tag{150}$$

and hence

$$\|X_r^\epsilon\|_F^2 = n_r^2 \tag{151}$$

exactly. Therefore, the within-block contribution to ϕ is constant under this perturbation.

Using the second-order expansion of ϕ, and noting that only cross-block terms vary to second order, we obtain

$$\phi(O^\epsilon) - \phi(O) = 2b\epsilon^2 n_r \Big(h^\top (X - X_r) h - \|h\|^2 \, v_r^\top (X - X_r) v_r \Big) + O(\epsilon^3). \tag{152}$$

where the factor n_r arises from summing contributions over nodes in B_r. Since O is a strict local maximizer modulo common rotations and $b > 0$, the quadratic form above must be nonpositive for all $h \perp v_r$. This inequality implies that v_r maximizes the Rayleigh quotient of $X - X_r$, and strict maximality ensures uniqueness. Therefore v_r is the unique top eigenvector of $X - X_r$.

Now observe that

$$X = (X - X_r) + n_r v_r v_r^\top. \tag{153}$$

On the orthogonal complement of v_r, the matrices X and $X - X_r$ coincide, while along v_r, the matrix X has the additional positive contribution n_r. Hence v_r is also the unique top eigenvector of X.

Applying this argument to each block yields

$$v_1 = \cdots = v_K = v, \tag{154}$$

and therefore

$$o_i \in \{\pm v\}, \qquad \forall i. \tag{155}$$

Thus, every strict local maximizer modulo common rotations of the clustered mean-field model is polarized.

G Regularity of the Recommendation-Augmented Dynamics

This section provides the details behind Lemma 3.1. The key idea is, although the two-step map is nonlinear, it is still regular enough for optimization, as it is Lipschitz and its Jacobian is Lipschitz on the feasible domain.

We summarize the two-step dynamics for fixed $O^{(t)}$, R, and W. Define

$$G := O^{(t)} R^\top, \tag{156}$$

$$Z(A) := O^{(t)} + (A \odot G) R, \tag{157}$$

$$\widehat{O}(A) := \mathcal{N}\big(Z(A)\big), \tag{158}$$

$$Y(X) := X + (W \odot X X^\top) X, \tag{159}$$

$$F(O^{(t)}; A) := \mathcal{N}\big(Y(\widehat{O}(A))\big). \tag{160}$$

The decomposition above separates the two-step dynamics into an affine recommendation update, a row normalization, a nonlinear social interaction map, and a final normalization. This decomposition is useful both conceptually and technically: each component has a clear meaning, and each one admits a clean regularity bound.

G.1 Row normalization

Lemma G.1. *Let $n(x) = x/\|x\|$ for $x \neq 0$. Then*

$$Dn(x) = \frac{1}{\|x\|}\left(I - \frac{xx^\top}{\|x\|^2}\right), \tag{161}$$

and

$$\|Dn(x)\|_{\mathrm{op}} \leq \frac{1}{\|x\|}, \qquad \|D^2 n(x)\|_{\mathrm{op}} \leq \frac{2}{\|x\|^2}. \tag{162}$$

Consequently, on the domain $\{x : \|x\| \geq 1\}$, the normalization map is 1-Lipschitz, and its Jacobian is 2-Lipschitz:

$$\|n(x) - n(y)\| \leq \|x - y\|, \qquad \|Dn(x) - Dn(y)\|_{\mathrm{op}} \leq 2\|x - y\|. \tag{163}$$

Proof. Differentiate $n(x) = x\|x\|^{-1}$ directly. The operator norm bounds follow from the explicit formula and the fact that $I - xx^\top/\|x\|^2$ is an orthogonal projector. For $\|x\|, \|y\| \geq 1$, the first inequality follows directly from

$$\|x - y\|^2 - \left\|\frac{x}{\|x\|} - \frac{y}{\|y\|}\right\|^2 = (\|x\| - \|y\|)^2 + 2(\|x\|\|y\| - 1)(1 - \cos\theta) \geq 0, \tag{164}$$

where θ is the angle between x and y. The Jacobian bound follows from the explicit expression for Dn, the reverse triangle inequality, and the corresponding bound between the normalized directions. □

Lemma G.2. *Every pre-normalized row in the recommendation step and the social interaction step has norm at least 1.*

Proof. Consider the recommendation step. The i-th pre-normalized row is

$$z_i = o_i^{(t)} + \sum_{k=1}^{M} a_{ik}\langle o_i^{(t)}, r_k\rangle r_k. \tag{165}$$

Expanding its squared norm gives

$$\|z_i\|^2 = \|o_i^{(t)}\|^2 + 2\sum_{k=1}^{M} a_{ik}\langle o_i^{(t)}, r_k\rangle^2 + \left\|\sum_{k=1}^{M} a_{ik}\langle o_i^{(t)}, r_k\rangle r_k\right\|^2. \tag{166}$$

Every term on the right-hand side is nonnegative and $\|o_i^{(t)}\| = 1$, so $\|z_i\| \geq 1$. The same argument applies to the social interaction step with r_k replaced by neighboring opinions and a_{ik} replaced by w_{ij}. □

Lemma G.2 ensures that the normalization maps remain away from their singularity on the feasible domain.

G.2 Proof of Regularity and Smoothness of the Updated Dynamics

Proof of Lemma 3.1. We prove the Lipschitz continuity of the full map by composing the Lipschitz bounds of its building blocks.

First, since $Z(A) = O^{(t)} + (A \odot G)R$, we have $Z(A) - Z(B) = ((A - B) \odot G)R$. Therefore

$$\|Z(A) - Z(B)\|_F \leq \|(A - B) \odot G\|_F \|R\|_2 \leq \|G\|_{\max}\|R\|_2\|A - B\|_F. \tag{167}$$

This shows that the recommendation pre-normalization map is Lipschitz. Intuitively, changing the exposure weights changes the additive recommendation term linearly, and the factor $\|R\|_2$ measures how strongly the recommendation geometry can amplify that change.

Next, by Lemmas G.1 and G.2 row normalization is 1-Lipschitz on the feasible domain, so

$$\|\widehat{O}(A) - \widehat{O}(B)\|_F \leq \|Z(A) - Z(B)\|_F. \tag{168}$$

Now consider

$$Y(X) = X + (W \odot XX^\top)X. \tag{169}$$

For two matrices X and P,

$$Y(X) - Y(P) = (X - P) + (W \odot XX^\top)(X - P) + \big(W \odot (XX^\top - PP^\top)\big)P. \tag{170}$$

Using

$$XX^\top - PP^\top = X(X - P)^\top + (X - P)P^\top, \tag{171}$$

and the fact that all rows remain unit-normalized after each normalization step, standard Frobenius/operator norm inequalities yield

$$\|Y(X) - Y(P)\|_F \le (2 + 2\|W\|_2)\|X - P\|_F. \tag{172}$$

The point of this estimate is that the nonlinear interaction map is still stable: its sensitivity is controlled by the spectral size of the network matrix.

Applying row normalization one last time, again using Lemmas G.1 and G.2 gives

$$\|F(O^{(t)}; A) - F(O^{(t)}; B)\|_F \le (2 + 2\|W\|_2)\|\widehat{O}(A) - \widehat{O}(B)\|_F. \tag{173}$$

Combining the previous bounds yields

$$\|F(O^{(t)}; A) - F(O^{(t)}; B)\|_F \le (2 + 2\|W\|_2)\|O^{(t)}R^\top\|_{\max}\|R\|_2\|A - B\|_F. \tag{174}$$

Under the standing unit-norm assumptions on the rows of $O^{(t)}$ and R, we have $\|O^{(t)}R^\top\|_\infty \le 1$. Since W is symmetric and row-stochastic, $\|W\|_2 = 1$. Therefore

$$\|F(O^{(t)}; A) - F(O^{(t)}; B)\|_F \le 4\|R\|_2\|A - B\|_F, \tag{175}$$

which proves the first part of Lemma 3.1.

The second part of the proof follows the same decomposition as the first part of the proof of Lemma 3.1, but now we control how its Jacobian varies with A.

The map $Z(A)$ is affine in A, so its derivative is constant and its second derivative is zero. Thus all curvature in the full dynamics comes from the two normalization operators and the social interaction map.

By Lemma G.1, row normalization has bounded first and second derivatives on the feasible domain. By Lemma G.1, row normalization is Lipschitz and has a Lipschitz-continuous Jacobian on the relevant domain.

For the interaction map

$$Y(X) = X + (W \odot XX^\top)X, \tag{176}$$

its Fréchet derivative in direction H is

$$DY_X[H] = H + (W \odot XX^\top)H + \big(W \odot (XH^\top + HX^\top)\big)X. \tag{177}$$

This expression is polynomial in X and linear in H, so the map $X \mapsto DY_X$ is Lipschitz on bounded sets. Since the rows of the normalized states remain on the sphere, we work on a bounded set automatically.

The composition rule for Jacobian-Lipschitz constants gives a bound of the form

$$\beta_{g\circ h} \le \beta_g L_h^2 + L_g\beta_h, \tag{178}$$

where L_h denotes a Lipschitz constant and β_h a Jacobian-Lipschitz constant. Applying this rule first to $\widehat{O}(A) = \mathcal{N}(Z(A))$ and then to $F(O^{(t)}; A) = \mathcal{N}(Y(\widehat{O}(A)))$, and substituting the previously established bounds, yields a valid global bound of the form

$$\|D_A F(O^{(t)}; A) - D_A F(O^{(t)}; B)\|_{\mathrm{op}} \le \beta_F\|A - B\|, \tag{179}$$

Now, we bound the mappings with smoothness constants,

$$L_Z \le \|R\|_2, \qquad \beta_Z = 0, \tag{180}$$

$$L_{\mathcal{N}} \le 1, \qquad \beta_{\mathcal{N}} \le 2. \tag{181}$$

Thus, for $\widehat{O}(A) = \mathcal{N}(Z(A))$,

$$L_{\widehat{O}} \le L_{\mathcal{N}} L_Z \le \|R\|_2, \tag{182}$$

and

$$\beta_{\widehat{O}} \le \beta_{\mathcal{N}} L_Z^2 + L_{\mathcal{N}} \beta_Z \le 2\|R\|_2^2. \tag{183}$$

For the social interaction map Y, on row-normalized inputs,

$$L_Y \le 2 + 2\|W\|_2 = 4, \qquad \beta_Y \le 6\|W\|_2 = 6. \tag{184}$$

Therefore, for $Y \circ \widehat{O}$,

$$L_{Y\circ\widehat{O}} \le L_Y L_{\widehat{O}} \le 4\|R\|_2, \tag{185}$$

and

$$\beta_{Y\circ\widehat{O}} \le \beta_Y L_{\widehat{O}}^2 + L_Y \beta_{\widehat{O}} \le 6\|R\|_2^2 + 2 \cdot 4\|R\|_2^2 = 14\|R\|_2^2. \tag{186}$$

Finally, applying the outer normalization gives

$$\beta_F \le \beta_{\mathcal{N}} L_{Y\circ\widehat{O}}^2 + L_{\mathcal{N}} \beta_{Y\circ\widehat{O}} \le 2(4\|R\|_2)^2 + 14\|R\|_2^2 = 46\|R\|_2^2. \tag{187}$$

under the unit-row and symmetric row-stochastic assumptions. This proves Lemma 3.1. □

H Proof of Objective Tractability Theorem

Theorem 3.1 shows that smoothness of the dynamics directly induces smoothness of a broad class of optimization objectives under the assumption that the metrics involved are also smooth. It establishes that measuring the dynamics with any smooth downstream metric preserves First-order regularity and results in a smooth objective function.

Proof of Theorem 3.1. Let the objective be

$$Q(A) = \sum_{i=1}^{P} \lambda_i \varphi_i(F(O^{(t)}; A)), \tag{188}$$

where each φ_i is continuously differentiable with Lipschitz-continuous gradient on the reachable set

$$\mathcal{R}_t := \{F(O^{(t)}; A) : A \in \mathcal{A}\}.$$

By Lemma 3.1, the map $A \mapsto F(O^{(t)}; A)$ is Lipschitz and has Lipschitz-continuous Jacobian on $\mathcal{A}$. So, applying the chain rule gives

$$\nabla Q(A) = \sum_{i=1}^{P} \lambda_i \, D_A F(O^{(t)}; A)^\top \nabla\varphi_i(F(O^{(t)}; A)). \tag{189}$$

To compare $\nabla Q(A)$ and $\nabla Q(B)$, we add and subtract the mixed term

$$D_A F(O^{(t)}; A)^\top \nabla\varphi_i(F(O^{(t)}; B))$$

for each i, yielding

$$\begin{aligned}\nabla Q(A) - \nabla Q(B) = \sum_{i=1}^{P} \lambda_i \Big[& D_A F(O^{(t)}; A)^\top \big(\nabla\varphi_i(F(O^{(t)}; A)) - \nabla\varphi_i(F(O^{(t)}; B))\big) \\ & + \big(D_A F(O^{(t)}; A)^\top - D_A F(O^{(t)}; B)^\top\big)\nabla\varphi_i(F(O^{(t)}; B))\Big].\end{aligned} \tag{190}$$

Since $\mathcal{A}$ is compact and F is continuous, the reachable set $\mathcal{R}_t$ is compact. Thus $D_A F$ and $\nabla\varphi_i$ are bounded on the relevant domains. Taking norms and using Lipschitz continuity of F, $D_A F$, and $\nabla\varphi_i$, we obtain

$$\|\nabla Q(A) - \nabla Q(B)\|_F \le L_Q \|A - B\|_F, \tag{191}$$

for some constant $L_Q > 0$, completing the proof. □

I Auxiliary Facts for the Zeroth-Order Method

This section provides details for Section 4. Since the optimization problem in Equation (18) is formulated as a minimization problem, the correct projected gradient mapping is the descent-form mapping associated with Euclidean projection onto the feasible set.

The off-feasible-set queries are also well-defined locally. By Lemma G.2, $\|z_i(A)\| \geq 1$ for every $A \in \mathcal{A}$, while $\|Z(A) - Z(B)\|_F \leq \|R\|_2 \|A - B\|_F$. Hence there exists $\rho > 0$ such that all recommendation pre-normalization rows remain bounded away from zero whenever $\mathrm{dist}(B, \mathcal{A}) < \rho$. The first normalization is therefore smooth on this neighborhood, and its output has unit-norm rows, so the social pre-normalization rows remain of norm at least 1. Consequently $F(O^{(t)}; B)$, and any smooth downstream objective $J(B)$, extend smoothly to this neighborhood. Since $U \sim \mathrm{Unif}(\mathbb{S}^{NM-1})$ satisfies $\|U\|_F = 1$, choosing $\mu < \rho$ guarantees that $A \pm \mu U$ lie in this neighborhood.

We therefore work with

$$\mathcal{G}_\eta^\mu(A) := \frac{1}{\eta}\Big(A - \Pi_{\mathcal{A}}(A - \eta \nabla J_\mu(A))\Big). \tag{192}$$

Lemma I.1. *Let*

$$\bar{A}^{t+1} := \Pi_{\mathcal{A}}\big(A^t - \eta \nabla J_\mu(A^t)\big). \tag{193}$$

Then

$$\langle \nabla J_\mu(A^t), A^t - \bar{A}^{t+1} \rangle \geq \frac{1}{\eta}\|A^t - \bar{A}^{t+1}\|_F^2 = \eta \|\mathcal{G}_\eta^\mu(A^t)\|_F^2. \tag{194}$$

Proof. The Euclidean projection onto a closed convex set is characterized by the variational inequality

$$\langle A^t - \eta \nabla J_\mu(A^t) - \bar{A}^{t+1}, A - \bar{A}^{t+1} \rangle \leq 0 \qquad \forall A \in \mathcal{A}. \tag{195}$$

Choosing $A = A^t$ and rearranging gives

$$\langle \nabla J_\mu(A^t), A^t - \bar{A}^{t+1} \rangle \geq \frac{1}{\eta}\|A^t - \bar{A}^{t+1}\|_F^2. \tag{196}$$

The identity with $\mathcal{G}_\eta^\mu$ follows from the definition of the projected gradient mapping. □

Lemma I.1 is the optimization analogue of the geometric decomposition in Section 3. It says that the projection step is not arbitrary. It still preserves enough alignment with the negative gradient to guarantee descent.

J Proof of Zeroth-order Convergence Corollary

Proof of Corollary 4.1. By Theorem 3.1 and standard properties of randomized smoothing, J_μ has an L_J-Lipschitz continuous gradient on $\mathcal{A}$. Moreover, by eq. (20),

$$\mathbb{E}[\widehat{g}_t \mid A^t] = \nabla J_\mu(A^t), \qquad \mathbb{E}\big[\|\widehat{g}_t - \nabla J_\mu(A^t)\|_F^2 \mid A^t\big] \leq \frac{\sigma^2}{q}. \tag{197}$$

Therefore, for $0 < \eta \leq 1/(2L_J)$, standard projected stochastic zeroth-order optimization results Ghadimi & Lan (2013); Liu et al. (2018) give

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\|\mathcal{G}_\eta^\mu(A^t)\|_F^2 \leq C_1 \frac{J_\mu(A^0) - \inf_{A \in \mathcal{A}} J_\mu(A)}{\eta T} + C_2 \eta \frac{\sigma^2}{q}, \tag{198}$$

for constants $C_1, C_2 > 0$. Since $\mathcal{A}$ is compact and J_μ is continuous, the initial optimality gap is finite. Absorbing this gap and the constants into the $O(\cdot)$ notation gives eq. (22). □

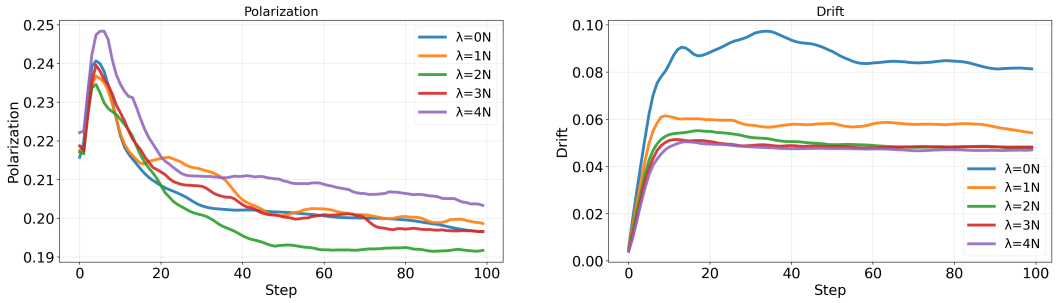

(a) Normalized squared alignment

(b) Drift

Figure 2: Sensitivity of regularization parameter λ. Increasing the penalty scale improves drift control, while normalized squared alignment has a non-monotonic relation with it.

K Transfer to the Original Objective

To transfer the stationarity bound from J_μ to J, define

$$\mathcal{G}_\eta(A) := \frac{1}{\eta}\Big(A - \Pi_{\mathcal{A}}(A - \eta\nabla J(A))\Big). \tag{199}$$

By non-expansiveness of the projection operator and the smoothing-bias bound,

$$\|\mathcal{G}_\eta(A) - \mathcal{G}^\mu_\eta(A)\|_F \le \|\nabla J(A) - \nabla J_\mu(A)\|_F \le B_\mu. \tag{200}$$

Therefore,

$$\|\mathcal{G}_\eta(A)\|_F^2 \le 2\|\mathcal{G}^\mu_\eta(A)\|_F^2 + 2B_\mu^2. \tag{201}$$

Averaging over t and applying corollary 4.1 gives

$$\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\|\mathcal{G}_\eta(A^t)\|_F^2 \le O\left(\frac{1}{\eta T}\right) + O\left(\eta\frac{\sigma^2}{q}\right) + O(B_\mu^2). \tag{202}$$

Thus, choosing $\eta = \Theta(T^{-1/2})$ yields the standard $O(T^{-1/2})$ average stationarity rate for fixed q, up to the smoothing-bias term.

L Additional Experiments

We present additional experiments to examine the behaviour of the dynamics, compare intervention methods, and test the sensitivity of the Zeroth-order optimizer. All experiments are conducted on an *Apple MacBook Air M4, 16GB RAM, 512GB SSD*. The Facebook dataset is obtained from the Stanford Network Analysis Project (SNAP) Leskovec & Mcauley (2012) and used in accordance with its BSD license and research usage guidelines. We use Facebook only as a real-world network topology; all opinion and recommendation states are synthetically generated.

L.1 Sensitivity Analysis: Regularization Parameter

We study the effect of the regularization parameter λ, which controls the trade-off between squared alignment and drift in the optimization objective. We test different values of λ on the same environment with 100 opinions and 5 recommendations using the Zeroth-order optimizer. We monitor normalized squared alignment and drift for 100 timesteps. For scaling between the metrics, the selected values of λ scale by N.

As seen in Figure 2, increasing λ places greater emphasis on controlling population drift. When $\lambda = 0$, drift is not explicitly penalized and is consequently larger and more variable. As λ increases, the observed drift becomes smaller and more stable.

The effect on normalized squared alignment is less monotonic. All tested values show a similar reduction pattern, while the moderate value $\lambda = 2N$ achieves the lowest alignment in this experiment. Thus, λ controls an empirical trade-off between reducing squared alignment and limiting population drift.

L.2 Uncontrolled Dynamics and Stronger Baselines

We additionally evaluate the model using noisy, community-correlated bipolar initializations. Unless otherwise stated, these experiments use $N = 100$, $M = 8$, $d = 10$, and four communities. SBM networks use within-community and cross-community edge probabilities 0.5 and 0.08, respectively. Opinions are initialized around two antipodal poles with Gaussian noise 0.22 before row normalization. Recommendations are constructed from a mixture of aligned, cross-cutting, diverse, and axis-reinforcing directions with noise 0.12. For the real-topology experiments, we sample connected 100-node subgraphs from the Facebook network and detect communities from the sampled topology.

Uncontrolled trajectories. For each predetermined seed, the intervention state is also rolled forward independently without recommendations until reaching a fixed-point residual below 10^{-3}, with a maximum of 500 steps. No seed is discarded post hoc. We monitor the first principal-direction share

$$\mathrm{PC1}(O) := \frac{\sigma_1(O)^2}{\sum_k \sigma_k(O)^2}, \tag{203}$$

the mean absolute loading on the dominant axis, and the fixed-point residual. We measure network disagreement as

$$\mathrm{Dis}(O) := \frac{1}{N} \sum_{i,j} w_{ij}\big(1 - \langle o_i, o_j \rangle\big).$$

Across five seeds, the terminal PC1 share is 1.0000 ± 0.0000 on SBM and 0.9993 ± 0.0007 on the real topology, while the corresponding fixed-point residuals are $(7.09 \pm 0.54) \times 10^{-4}$ and $(9.52 \pm 0.09) \times 10^{-4}$, respectively. The mean terminal axial concentrations are 1.0000 ± 0.0000 and 0.9996 ± 0.0003. The trajectories require 18.6 ± 1.5 and 81.0 ± 10.6 steps, respectively. These experiments show that the proposed dynamics can evolve from noisy bipolar community structure to strongly axially concentrated approximate fixed points. In the corresponding bipolar experiments below, this strong directional concentration coexists with substantial disagreement, supporting the interpretation of high squared alignment as polarization rather than purely consensus-like concentration. They complement, rather than replace, the non-polarized fixed-point examples in Appendix F.2.

Baseline definitions. We use uniform exposure as the feasible, action-matched reference. The *Similarity* baseline assigns each agent the recommendation with maximum current cosine similarity. The *Best-of-K random* baseline samples row-wise simplex exposures from a uniform Dirichlet distribution and retains the lowest-objective sample, using the same total objective-evaluation budget as the Zeroth-order optimizer.

The *PCA diversity* baseline first extracts the top two principal directions of the current opinion matrix. For every agent i and recommendation r_k, it computes the recommendation-only candidate

$$\widetilde{o}_{ik} = \mathcal{N}(o_i + \langle o_i, r_k \rangle r_k)\,, \tag{204}$$

and assigns the recommendation whose candidate has the smallest squared projection onto the current dominant two-dimensional subspace.

The topology heuristic uses only graph structure. Let $d_G(i, j)$ be undirected shortest-path distance and define

$$P_{ij} = \begin{cases} \big(d_G(i,j) + \varepsilon\big)^{-1}, & i \neq j, \\ 0, & i = j. \end{cases} \tag{205}$$

Nodes are processed in decreasing order of $\sum_j P_{ij}$. If c_m denotes the number of previously assigned nodes receiving recommendation m, node i is assigned the recommendation minimizing

$$\sum_{\substack{j \text{ assigned} \\ m_j = m}} P_{ij} + 0.1 \left(c_m - \frac{N}{M} \right)^2. \tag{206}$$

Ties are resolved in favor of the currently least-used recommendation. Thus, nearby nodes are encouraged to receive different recommendations while maintaining approximately balanced recommendation counts.

Table 1: Stronger baseline comparison. Objective and alignment reductions are measured relative to uniform exposure; disagreement is reported directly.

Topology	Method	Obj. red. (%)	Align. red. (%)	Disagreement	Drift	Runtime (s)
SBM	Uniform	0.00 ± 0.00	0.00 ± 0.00	0.4488 ± 0.0126	0.0851 ± 0.0100	0.000 ± 0.000
	Best-of-K random	2.43 ± 0.32	1.74 ± 0.30	0.4539 ± 0.0120	0.0769 ± 0.0100	2.707 ± 0.070
	Similarity	-2.13 ± 0.92	-1.12 ± 0.61	0.4438 ± 0.0127	0.0959 ± 0.0099	0.000 ± 0.000
	PCA diversity	35.12 ± 1.52	30.17 ± 1.62	0.5460 ± 0.0071	0.0107 ± 0.0018	0.000 ± 0.000
	Topology heuristic	3.68 ± 1.20	2.97 ± 0.91	0.4580 ± 0.0112	0.0756 ± 0.0094	0.005 ± 0.000
	Network-aware ZO	16.04 ± 1.40	12.70 ± 1.54	0.4879 ± 0.0100	0.0415 ± 0.0068	2.685 ± 0.010
Real topology	Uniform	0.00 ± 0.00	0.00 ± 0.00	0.2122 ± 0.1018	0.0746 ± 0.0067	0.000 ± 0.000
	Best-of-K random	2.21 ± 0.38	1.55 ± 0.33	0.2184 ± 0.1019	0.0670 ± 0.0071	2.700 ± 0.030
	Similarity	-2.47 ± 1.07	-1.36 ± 0.61	0.2064 ± 0.1004	0.0865 ± 0.0048	0.000 ± 0.000
	PCA diversity	30.47 ± 1.53	25.97 ± 1.72	0.3285 ± 0.0916	0.0107 ± 0.0024	0.000 ± 0.000
	Topology heuristic	4.24 ± 2.77	3.20 ± 2.17	0.2259 ± 0.1052	0.0621 ± 0.0128	0.005 ± 0.000
	Network-aware ZO	13.16 ± 2.23	10.33 ± 2.06	0.2554 ± 0.0995	0.0391 ± 0.0089	2.707 ± 0.057

Finally, *Network-aware ZO* optimizes the full one-step recommendation and social-interaction objective. These experiments use $T = 50$, $q = 64$, $\mu = 0.01$, step size 0.05, and $\lambda = N$. Results are reported as mean $\pm$ standard deviation over five seeds.

The network-aware Zeroth-order method consistently improves over uniform exposure, budget-matched random search, similarity-based exposure, and the topology-only heuristic on both network types. Importantly, the experiments retain substantial disagreement alongside strong directional concentration: under uniform exposure, disagreement is 0.4488 ± 0.0126 on SBM and 0.2122 ± 0.1018 on the real topology. Thus, the high squared-alignment states in these bipolar settings are not simply consensus-like configurations. The PCA-diversity heuristic performs particularly strongly in this bipolar setting, showing that direct access to an informative low-dimensional opinion geometry can provide a strong specialized baseline. These results therefore provide a stronger comparison than topology-only intervention while also illustrating the value of geometric information.

L.3 Zeroth-Order Sensitivity and Computational Cost

We next vary the smoothing radius μ, number of perturbation directions q, number of recommendations M, opinion dimension d, and number of agents N. Unless the corresponding quantity is varied, the base setting uses $N = 100$, $M = 8$, $d = 10$, $q = 64$, $\mu = 0.01$, and $T = 50$. Each sensitivity result is averaged over three seeds. Uniform exposure is again used as the action-matched reference.

Performance is nearly unchanged across the tested smoothing radii. Increasing q produces progressively larger objective and alignment reductions, at the cost of proportionally more objective evaluations and runtime, illustrating the expected variance–evaluation trade-off. Across the remaining sweeps, the optimizer improves over uniform exposure for every tested value of M, d, and N. With a fixed optimization budget, the magnitude of the improvement decreases as M or N increases, while runtime grows most substantially with N.

For T iterations and q two-point directions, the gradient estimator requires $2qT$ objective evaluations. The experimental implementation also evaluates the initial exposure once and the projected candidate once per iteration, giving

$$1 + T(2q + 1) \tag{207}$$

total optimization calls. Thus, for $T = 50$, the tested values $q = 16, 32, 64, 128$ require 1651, 3251, 6451, and 12851 objective calls, respectively. The Best-of-K random baseline above uses $K = 6451$, matching the total evaluation budget of the default $q = 64$ Zeroth-order run.

Implementation convention. The theoretical development in Section 4 uses bounded uniform-sphere directions to state the randomized-smoothing guarantee. The experiments reported in this section were generated with the Gaussian two-point implementation used in our experimental code. We retain these results unchanged; both implementations use symmetric two-point objective queries followed by projection onto $\mathcal{A}$.

Table 2: Sensitivity and scaling of the Zeroth-order method. Percentage reductions are measured relative to uniform exposure.

Parameter	Value	Obj. red. (%)	Align. red. (%)	Drift	Runtime (s)	Obj. calls
μ	0.001	15.68 ± 0.90	12.41 ± 0.82	0.0430 ± 0.0022	2.684 ± 0.010	6451
	0.01	15.68 ± 0.91	12.41 ± 0.82	0.0430 ± 0.0022	2.665 ± 0.006	6451
	0.05	15.66 ± 0.93	12.39 ± 0.84	0.0430 ± 0.0023	2.684 ± 0.023	6451
	0.1	15.57 ± 1.00	12.33 ± 0.85	0.0434 ± 0.0026	2.739 ± 0.046	6451
q	16	8.58 ± 0.32	6.57 ± 0.42	0.0604 ± 0.0029	0.697 ± 0.012	1651
	32	11.82 ± 1.08	9.24 ± 0.92	0.0525 ± 0.0015	1.360 ± 0.005	3251
	64	15.68 ± 0.91	12.41 ± 0.82	0.0430 ± 0.0022	2.665 ± 0.004	6451
	128	20.63 ± 1.30	16.71 ± 1.11	0.0329 ± 0.0013	5.269 ± 0.021	12851
M	4	19.85 ± 1.70	15.94 ± 1.36	0.0351 ± 0.0032	2.652 ± 0.009	6451
	8	15.68 ± 0.91	12.41 ± 0.82	0.0430 ± 0.0022	2.703 ± 0.034	6451
	16	9.43 ± 0.06	7.28 ± 0.12	0.0580 ± 0.0015	2.742 ± 0.070	6451
	32	5.75 ± 0.26	4.27 ± 0.38	0.0669 ± 0.0003	2.728 ± 0.002	6451
d	2	3.96 ± 2.87	2.98 ± 2.66	0.0090 ± 0.0025	2.656 ± 0.012	6451
	5	10.57 ± 1.71	8.00 ± 1.56	0.0328 ± 0.0021	2.683 ± 0.006	6451
	10	15.68 ± 0.91	12.41 ± 0.82	0.0430 ± 0.0022	2.750 ± 0.023	6451
	20	22.14 ± 1.16	19.34 ± 1.42	0.0339 ± 0.0078	2.834 ± 0.020	6451
N	50	19.09 ± 1.27	15.12 ± 1.29	0.0381 ± 0.0041	1.386 ± 0.007	6451
	100	15.68 ± 0.91	12.41 ± 0.82	0.0430 ± 0.0022	2.752 ± 0.001	6451
	200	11.48 ± 0.34	9.01 ± 0.23	0.0514 ± 0.0019	6.964 ± 0.065	6451
	400	8.11 ± 0.38	6.30 ± 0.29	0.0607 ± 0.0009	23.660 ± 0.172	6451

L.4 Sensitivity Analysis: Number of Nodes and Network Density

To understand sensitivity to network parameters, we additionally conduct the original experiments with varying numbers of opinions and network density. We keep the number of recommendations fixed to $M = 10$ and the number of communities to 5. The synthetic networks are sampled from a Stochastic Block Model (SBM). Network density is controlled by **inter_p** and **intra_p**, which respectively determine cross-community and within-community edge probabilities. The results are shown in Figure 3.

Across the displayed settings, the uncontrolled dynamics attain high normalized squared alignment. This indicates strong directional concentration, but squared alignment alone does not distinguish consensus from bipolar polarization. The topology heuristic reduces squared alignment relative to no recommendation, although it does not use the opinion geometry. Both optimization methods consistently outperform the topology heuristic in these experiments. The First-order method performs best, while the Zeroth-order method remains close in performance across the tested network sizes and densities.

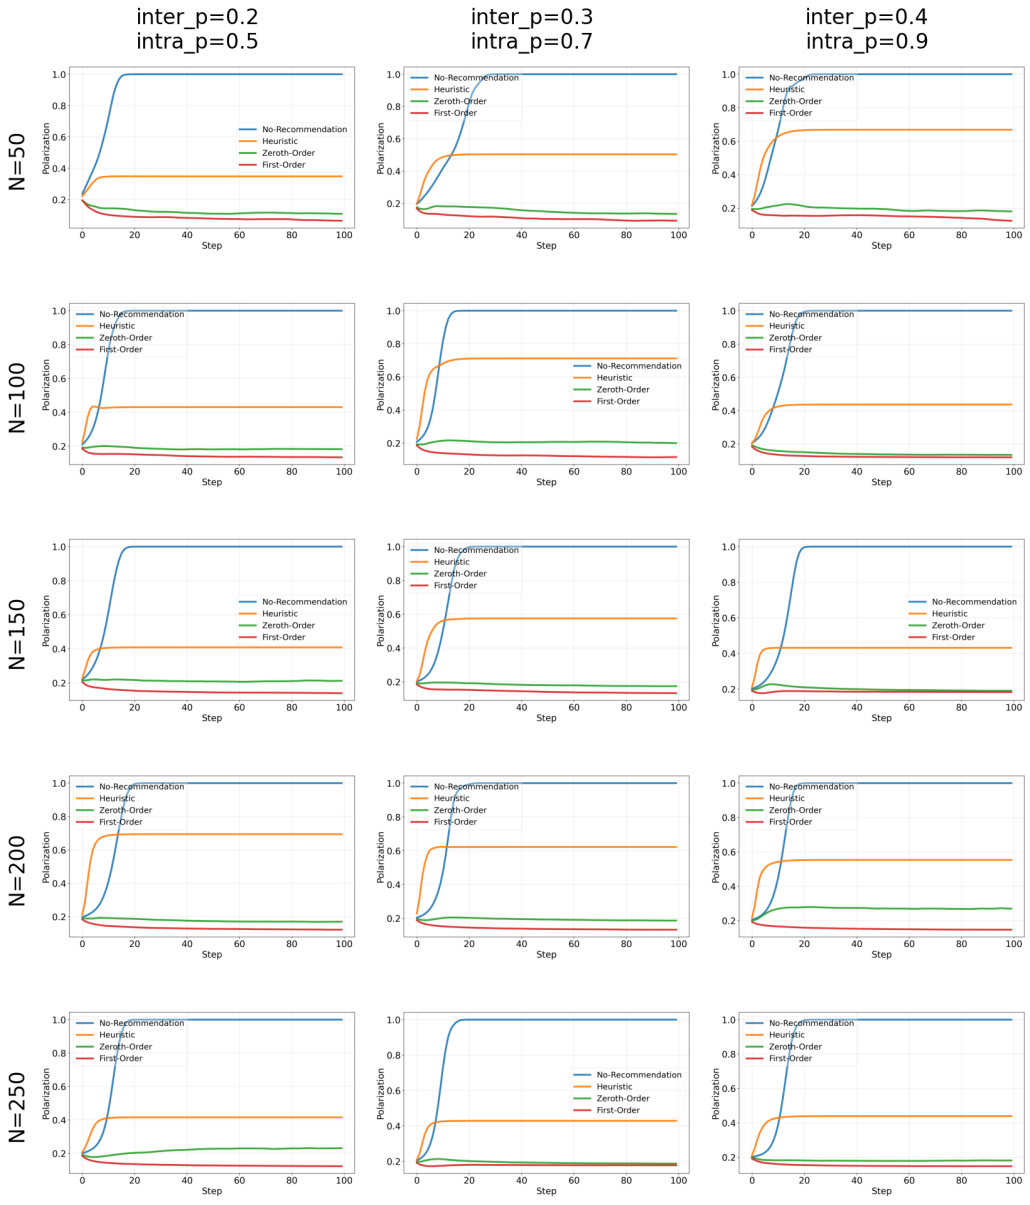

Figure 3: Sensitivity analysis for varying number of opinions and network density.