

OPINION DYNAMICS ON DISCOURSE SHEAVES*

JAKOB HANSEN[†] AND ROBERT GHRIST[‡]

Abstract. We introduce a novel class of Laplacians and diffusion dynamics on *discourse sheaves* as a model for network dynamics, with application to opinion dynamics on social networks. These *sheaves* are algebraic data structures tethered to a network (or more general space) that can represent various modes of communication, including selective opinion modulation and lying. After introducing the sheaf model, we develop a sheaf Laplacian in this context and show how to evolve both opinions and communications with diffusion dynamics over the network. Issues of controllability, reachability, bounded confidence, and harmonic extension are addressed using this framework.

Key words. social network, cellular sheaves, opinion dynamics, Laplacians

AMS subject classifications. 91D30, 55N30 (MSC2020)

1. Introduction. Social networks are one of the principal motivating examples for the study of complex networks. Among the many interesting problems associated with social networks, *opinion dynamics* — the study of how preferences or opinions emerge and evolve — are especially interesting, blending ideas from dynamical systems and graph theory. Structural effects of the network on opinion dynamics began with the analysis of linear dynamical models [36, 15, 20] and have developed into more sophisticated formulations [14, 17, 25], including features such as bounded confidence. This paper introduces both a novel model and a novel set of tools for the analysis of opinion dynamics. After a brief review of classical opinion dynamics models, we survey the results of this paper.

1.1. Models of Opinion Dynamics. Early models of opinion dynamics used linear network dynamics to evolve single-dimensional preferences. Consider a social network represented as an undirected graph $G = (V, E)$ of vertices and edges. The state space for single-opinion real dynamics is \mathbb{R}^V , with $x \in \mathbb{R}^V$ representing a distribution of preferences $x_v \in \mathbb{R}$ at each vertex, ranging from positive to indifferent (null) to negative. Graph-based linear dynamics evolve preferences over time.

In continuous-time models, the graph Laplacian, L , generates dynamics via the graph diffusion equation

$$(1.1) \quad \frac{dx}{dt} = -\alpha Lx, \quad : \quad \alpha > 0,$$

perhaps with slight modifications [1, 36]. Analogous models [18, 15, 29] were studied in discrete time, with a (typically stochastic) state evolution matrix following the sparsity pattern of the adjacency matrix, A , of the social network:

$$(1.2) \quad x[t + 1] = Ax[t].$$

Variations with more terms added additional richness to the models [20].

*Submitted to the editors DATE.

Funding: This work was funded by the Office of the Assistant Secretary of Defense Research & Engineering through a Vannevar Bush Faculty Fellowship, ONR N00014-16-1-2010.

[†]Department of Mathematics, University of Pennsylvania, Philadelphia, PA (jhansen@math.upenn.edu)

[‡]Department of Mathematics and Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA (ghrist@math.upenn.edu).

Without modification, these linear graph models result in asymptotically stable equilibria at a single consensus opinion over the network. While this global fixed consensus may be useful in some situations (e.g., flocking in robotics/swarm applications [26, 35]), it does not represent the typical behavior of opinion distributions in social networks. Indeed, a central problem in the study of opinion dynamics is the construction of simple models that replicate one of the most salient features of real-world opinions: the existence of polarization or failure to come to consensus, known as the *community cleavage problem* [19]. The earliest approaches to solving this problem added extra constant terms to dynamics of the form (1.1) or (1.2) that encouraged diversity in opinions across the network [36, 20, 21]. These extra terms were interpreted as external influences or the effect of individual stubbornness on opinions.

Much recent work has focused on a nonlinear extension of the discrete-time dynamics that adds *bounded confidence* to agents' evaluation of their neighbors' opinions [14, 17, 25]. The most famous such model, popularized by Hegselmann and Krause, posits a threshold r for opinion sharing. If agents i and j have opinions that differ by more than r , they do not communicate. Otherwise, they influence each other's opinions linearly as in the discrete-time dynamics above. That is,

$$(1.3) \quad x_i[t+1] = \sum_{i \sim j} A_{ij} x_j \mathbf{1}_{|x_i - x_j| < r}.$$

These dynamics are complex enough to produce bifurcated opinion distributions without external influences, while admitting some direct analysis.

Not all work on opinion dynamics has used the Laplacian or adjacency-based linear network interaction formulations above. Other popular models include dynamics with discrete opinion spaces [34, 8] based on formalisms from statistical physics, as well as non-agent-based models that seek to understand the overall distribution of opinions in a population without tracking any individual's stance [37, 7].

In this paper, we will focus on network-based continuous-time models for opinion dynamics with continuous opinion spaces; discrete-time examples can typically be extracted from these via Euler discretization. Our models will find most kinship with those that approach opinion dynamics from a control-theoretic or systems analysis perspective, like those discussed in [31, 32].

1.2. Contributions of This Work. In this paper, we introduce a novel approach to networked opinion dynamics, using ideas from *sheaf theory*. This subject, commonly used in algebraic topology and algebraic geometry, is vast [27, 16, 22, 24], but has a simple, graph-based reduction that amounts to little more than a networked system of linear transformations (see §2 for definitions). There are three key ideas from sheaf theory that we use:

1. *Cellular sheaves* are a topological data structure for graphs (or more general cell complexes) [13], which, we argue, permits modeling of very sophisticated opinion state spaces and communication strategies.
2. *Sheaf cohomology* is an algebraic-topological invariant of the sheaves over graphs. We demonstrate its use in computing obstructions to solving problems of opinion dynamics on graphs.
3. *Sheaf Laplacians* are far-reaching extensions of the graph Laplacian [23] and extend ideas of harmonic flow to a wide array of opinion dynamics models.

Although not a familiar toolset within network science, sheaf-theoretic methods have a remarkable range of expressiveness and can encode more realistic modes of expression of opinions, including exaggeration, modulation, and selective obfuscation.

After setting up the relevant mathematical structures in §2, we proceed directly to the contributions, summarized below.

1. In §3, we introduce our model of opinions over a social network via a cellular sheaf. In this model, each agent has an *opinion space* (a vector space with dedicated basis) and each edge has an independent *discourse space* (a vector space with basis of topics up for discussion). Expression of opinions is programmed via linear transformations from opinion to discourse spaces, allowing for private opinions selectively expressed or combined into policies.
2. In §4, we use the sheaf Laplacian to set up diffusion dynamics on the discourse sheaf, proving asymptotic convergence of initial opinions to a (literal) *harmonic* state: all agents express opinions in harmony with neighbors, though the private opinions of neighbors may be distinct (or even incomparable).
3. In §5, we use sheaf cohomology to characterize whether certain problems of extension and convergence have solutions. For example, if certain agents are inflexible and will not modify their opinions, does there exist a unique global solution through modification of others' opinions? We show that this is a problem of *harmonic extension*, determinable via a cohomology computation.
4. The question of manipulation of a system through inflexible agents leads naturally to questions of a control-theoretic nature. In §6, sheaf cohomology is shown to determine the controllability and observability of opinions.
5. Using the Laplacian to evolve individual opinions in order to come to harmonic expression is only half the picture. One could instead keep opinions fixed and evolve the expression of opinions in order to reduce discord. In §8, we extend the Laplacian diffusion model to the sheaf maps that express opinions. This leads to the interesting phenomenon of agents “learning to lie” to reach concord. The natural extension to joint opinion-expression diffusion is given in §9.
6. Finally, in §10-12, we move from linear to (slightly) nonlinear dynamics of opinion distributions on discourse sheaves, showing how to mimic the *bounded confidence* models of [25] and the antagonistic social dynamics of [4].

2. Introduction to Cellular Sheaves. For purposes of this paper, a cellular sheaf is a data structure augmenting a graph that describes consistency relationships for algebraic data attached to the graph. For simplicity, we work with vector spaces and linear transformations thereof. At one or two points, it will be convenient if not necessary to have all vector spaces real, finite-dimensional, and with an inner product structure. For the remainder of this work, we will implicitly assume these conditions: see [23, Section 3.3] for details on how to deal with more general Hilbert spaces. Further, for simplicity, we will assume that every vector space has a canonical orthonormal basis identifying it with \mathbb{R}^n for some n , and identify linear maps with their representing matrices and the inner product $\langle x, y \rangle$ with the standard inner product $x^T y$ on \mathbb{R}^n .

2.1. Definitions. Let G be a graph. A cellular sheaf \mathcal{F} on G is specified by the following data:

- a vector space $\mathcal{F}(v)$ for each vertex v of G ,
- a vector space $\mathcal{F}(e)$ for each edge e of G , and
- a linear map $\mathcal{F}_{v \triangleleft e} : \mathcal{F}(v) \rightarrow \mathcal{F}(e)$ for each incident vertex-edge pair $v \triangleleft e$.

The vector spaces $\mathcal{F}(v)$ are called the *stalks* over v , and the linear maps $\mathcal{F}_{v \triangleleft e}$ are the *restriction maps*. The terminology seems unmotivated and is inherited from more general sheaf theory [13]. Thinking in terms of data (stalks) and communication




FIG. 1. A cartoon diagram of a cellular sheaf over a graph. Vector spaces of varying dimensions over vertices and edges are attached via linear transformations, following the adjacency pattern of the graph. The entire system of linear transformations forms the cellular sheaf.

(restriction maps) is perhaps preferable in the context of this paper, see Figure 1.

The simplest example of a cellular sheaf is the *constant sheaf*. For V a fixed vector space, the constant V -sheaf on G , denoted \underline{V} , attaches a copy of V to each vertex and edge (all stalks are V), with $\underline{V}_{v \triangleleft e} = \text{id}_V$ for all incident vertex-edge pairs (all restriction maps are the identity). One interprets a constant sheaf as specifying that all vertices access the same data (from V) and communicate perfectly with their neighbors over edges.

There is more to a sheaf than its stalks: replacing all the identity maps in \underline{V} with zero-maps ($\mathcal{F}_{v \triangleleft e} : V \mapsto 0$) gives a very different sheaf in which all communication is devoid of information. Such a sheaf is a jumble in contrast to the tightly coordinated constant sheaf \underline{V} .

2.2. Sections and Sheaf Cohomology. If one visualizes a sheaf of vector spaces over a graph as being a network of vector spaces and linear transformations, then one is naturally led to questions of how to generalize the familiar notions of linear algebra — kernels, images, etc. — to such a networked structure. This is the impetus for *homological algebra* and the *cohomology* of a sheaf.

One begins by bundling all the data over vertices and over edges into a pair of conglomerated vector spaces. These are called spaces of *cochains*

$$(2.1) \quad C^0(G; \mathcal{F}) = \bigoplus_{v \in V(G)} \mathcal{F}(v)$$

$$(2.2) \quad C^1(G; \mathcal{F}) = \bigoplus_{e \in E(G)} \mathcal{F}(e).$$

Elements of C^0 are called *0-cochains*: such an $x \in C^0(G; \mathcal{F})$ consists of a choice of data, $x_v \in \mathcal{F}(v)$, for every vertex v of G . Likewise, an element $y \in C^1(G; \mathcal{F})$ is called a *1-cochain* and is a choice of data indexed over the edges of G .

In the same manner that edges and vertices are stitched together to form a graph, data over the vertices (0-cochains) and edges (1-cochains) are tied together via a linear transformation — the *coboundary* map, $\delta : C^0(G; \mathcal{F}) \rightarrow C^1(G; \mathcal{F})$. To define δ explicitly, choose a fixed but arbitrary orientation on each edge e . Then the evaluation of δ on an oriented edge $e = u \rightarrow v$ is defined as follows:

$$(2.3) \quad (\delta x)_e = \mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u.$$

The orientation merely serves as a choice of basis elements for defining the difference operation in δ . That choice is irrelevant, as what one cares about is the kernel of δ .

If one thinks about the coboundary δ as a measure of “disagreement” of data across an edge, then $\ker \delta$ is the subspace of $C^0(G; \mathcal{F})$ consisting of choices of data over the vertex set which “agree” over the edges. That is, for every $e = u \sim v$, $\mathcal{F}_{u \triangleleft e} x_u = \mathcal{F}_{v \triangleleft e} x_v$. The set of all such solutions to the global constraint satisfaction problem of the sheaf has the structure of a vector subspace of $C^0(G; \mathcal{F})$ and goes by the title of *cohomology*.

For \mathcal{F} a cellular sheaf on a graph G , the *zeroth cohomology* $H^0(G; \mathcal{F})$, also known as the space of *global sections* of \mathcal{F} , is

$$(2.4) \quad H^0(G; \mathcal{F}) = \ker \delta \subset C^0(G; \mathcal{F}).$$

The global sections of a sheaf are the global solutions to the networked system of constraint equations programmed into the restriction maps $\mathcal{F}_{u \triangleleft e}$ of the data structure over the graph. The cohomological terminology comes from the algebraic topology of sheaves [27, 16, 22] — a beautiful theory that can be ignored by the end-user of the models in this paper (but which secretly animates many of the results).

Related to the space of global sections are the subspaces of *local sections*. For A a subgraph of G , we let $C^0(A; \mathcal{F}) = \bigoplus_{v \in V(A)} \mathcal{F}(v)$ and $C^1(A; \mathcal{F}) = \bigoplus_{e \in E(A)} \mathcal{F}(e)$. The coboundary δ restricts to a map $C^0(A; \mathcal{F}) \rightarrow C^1(A; \mathcal{F})$; its kernel is the space of local sections over A , denoted $H^0(A; \mathcal{F})$. This is the subspace of $C^0(A; \mathcal{F})$ which is consistent over every edge in A .

Dual to the local sections over A is the cohomology relative to A . We let $C^0(G, A; \mathcal{F}) = \bigoplus_{v \notin V(A)} \mathcal{F}(v)$ and $C^1(G, A; \mathcal{F}) = \bigoplus_{e \notin E(A)} \mathcal{F}(e)$, and restrict δ to a map between these spaces. The *degree 0 relative cohomology* is $H^0(G, A; \mathcal{F}) = \ker \delta|_{C^0(G, A; \mathcal{F})}$. Local sections $H^0(A; \mathcal{F})$ should be thought of as assignments to a subset of vertices that are consistent on A , while the relative cohomology $H^0(G, A; \mathcal{F})$ can be viewed as global sections of \mathcal{F} on G that vanish on A .

Example 2.1. For the constant sheaf \mathbb{R} on a connected graph G , $H^0(G; \mathbb{R})$ is a one-dimensional vector space spanned by the constant functions on vertices of G . For any nonempty subgraph A of G , $H^0(A; \mathbb{R})$ consists of functions on the vertices of A which are locally constant on the subgraph; its dimension is the number of connected components of A . The relative cohomology $H^0(G, A; \mathbb{R})$ is zero-dimensional, since any constant function which is zero on A must be zero everywhere on G .

2.3. The Sheaf Laplacian. Recall that for a graph with signed incidence matrix B , the *graph Laplacian* is given by $L = BB^T$. These are combinatorial versions of the familiar second-order differential operator with enormous applicability across combinatorics, data science, and more [11, 5, 12]. Less well-known in applied mathematics is the definition of Laplacians of complexes of sheaves of inner-product spaces over topological spaces: these arise in Hodge theory and algebraic geometry [30]. Between the graph Laplacian and Hodge Laplacian lies a mean notion of a Laplacian for cellular sheaves on graphs [23].

The construction is uncomplicated. Observe that the coboundary δ of the constant sheaf \mathbb{R} with stalk \mathbb{R} equals the transposed signed incidence matrix B^T of the graph G . For a sheaf \mathcal{F} over G , δ may be seen as a generalized incidence matrix for \mathcal{F} . The potential variation in stalk dimensions means that δ is a block matrix, and the restriction maps determine the block entries, with sparsity pattern determined by the structure of G . For a sheaf \mathcal{F} on a graph G , the *sheaf Laplacian* is

$$(2.5) \quad L_{\mathcal{F}} = \delta^T \delta: C^0(G; \mathcal{F}) \rightarrow C^0(G; \mathcal{F}).$$

Just as the graph Laplacian does not depend on the orientations chosen for the edges in the signed incidence matrix, the sheaf Laplacian does not depend on the choice of orientations for the construction of the coboundary δ .

The following theorem may seem trivial in the context of sheaves over graphs, but it stems from deeper results (on sheaves over higher-dimensional cell complexes and with relations to higher cohomologies).

THEOREM 2.2 (Hodge Theorem). *For \mathcal{F} a sheaf on a graph G as above,*

$$(2.6) \quad H^0(G; \mathcal{F}) = \ker L_{\mathcal{F}}.$$

Example 2.3. The well-known fact that the kernel of the graph Laplacian is the space of locally constant functions on G follows from applying the Hodge theorem to the constant sheaf on G .

A straightforward computation shows that for a 0-cochain $x \in C^0(G; \mathcal{F})$, the value of the sheaf Laplacian at a given vertex v is

$$(2.7) \quad (L_{\mathcal{F}}x)_v = \sum_{v, u \triangleleft e} \mathcal{F}_{v \triangleleft e}^T (\mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u).$$

This implies that the matrix of $L_{\mathcal{F}}$ has a block structure, with diagonal blocks $L_{vv} = \sum_{v \triangleleft e} \mathcal{F}_{v \triangleleft e}^T \mathcal{F}_{v \triangleleft e}$ and off-diagonal blocks $L_{vu} = -\mathcal{F}_{v \triangleleft e}^T \mathcal{F}_{u \triangleleft e}$.

In the next section, we will interpret $(L_{\mathcal{F}}x)_v$ as a measure of the average disagreement of agent v with its neighbors, or equivalently, of the disagreement of v with an average of its neighbors.

Example 2.4. The sheaf in Figure 2 has a coboundary map represented by the matrix

$$(2.8) \quad \delta = \left[\begin{array}{c|ccc|ccc} -2 & 1 & -2 & 0 & 0 & 0 \\ \hline 0 & 1 & 1 & -1 & 0 & 0 \\ \hline 0 & 0 & 0 & -1 & -1 & 1 \\ \hline 0 & 0 & 0 & 0 & 0 & 1 \\ \hline 1 & 0 & 0 & 0 & -1 & 0 \end{array} \right],$$

and therefore its Laplacian is

$$(2.9) \quad L_{\mathcal{F}} = \delta^T \delta = \left[\begin{array}{c|ccc|ccc} 5 & -2 & 4 & 0 & -1 & 0 \\ \hline -2 & 2 & -1 & -1 & 0 & 0 \\ \hline 4 & -1 & 5 & -1 & 0 & 0 \\ \hline 0 & -1 & -1 & 2 & 1 & -1 \\ \hline -1 & 0 & 0 & 1 & 2 & -1 \\ \hline 0 & 0 & 0 & -1 & -1 & 2 \end{array} \right].$$

The block columns of δ and $L_{\mathcal{F}}$ correspond to the stalks $\mathcal{F}(v_i)$ for $i = 1, \dots, 4$.

3. Discourse Sheaves. We introduce a cellular sheaf model for opinions and discourse for which Laplacian diffusion is an effective and computable method.

Construction: Given a social network presented as a graph G with vertices representing agents and (undirected) edges representing pairwise communication, consider the following *discourse sheaf* \mathcal{F} . Each agent (vertex) v has an *opinion space*, a real vector space with basis some collection of topics. As with classical opinion dynamics models, points on each axis correspond to negative, neutral, or positive opinions




FIG. 2. A cellular sheaf over a simple cyclic graph. All stalks are of dimension one or two.




FIG. 3. In a discourse sheaf, stalks over vertices are individual opinion spaces, stalks over edges are discourse spaces, and restriction maps are expressions of opinions on the topics of discourse, formulated linearly from basis opinions.

on the topic, with a positive/negative intensity registered by the scalar value. This opinion space comprises the stalk $\mathcal{F}(v)$ of the discourse sheaf over v . A choice of element $x_v \in \mathcal{F}(v)$ is a vector recording the intensities of opinions or preferences on each of the *basis topics*.

Given an edge e between vertices u and v , it is presumed that there is a certain set of basis topics about which the two agent discuss. These are not necessarily the same as any of the basis topics from which $\mathcal{F}(u)$ or $\mathcal{F}(v)$ are generated; however, they do form the basis of an abstract *discourse space*, $\mathcal{F}(e)$, the stalk over e .

Each agent represents their opinions on the topics of discussion by formulating stances as a linear combination of existing opinions on personal basis topics. These expressions of opinion are linear transformations $\mathcal{F}_{u \triangleleft e} : \mathcal{F}(u) \rightarrow \mathcal{F}(e)$ and $\mathcal{F}_{v \triangleleft e} : \mathcal{F}(v) \rightarrow \mathcal{F}(e)$. If the agents hold opinions $x_u \in \mathcal{F}(u)$ and $x_v \in \mathcal{F}(v)$, then they have expressed consensus when $\mathcal{F}_{u \triangleleft e}(x_u) = \mathcal{F}_{v \triangleleft e}(x_v)$. This is the local consistency condition implicit in the construction of a sheaf — each edge imposes a linear consistency constraint on the stalks of its incident vertices. Note that this does not imply that u and v have the same opinions: it means that their expressions of personally held opinions have the appearance of agreement.

The apparatus of §2 becomes clearer in the context of discourse sheaves.

- The constant sheaf \mathbb{R}^n is a discourse sheaf in which every agent has an opinion on the same n basis topics, all of which are precisely expressed and discussed without embellishment. This is the implicit structure in most of the literature on opinion dynamics.
- A 0-cochain $x \in C^0(G; \mathcal{F})$ is a private opinion distribution on the set of agents.
- A 1-cochain $\xi \in C^1(G; \mathcal{F})$ is a distribution of expressed opinions over all pairwise agent discussions in the network.

- The coboundary map $\delta : C^0 \rightarrow C^1$ registers the “difference of expression” between agents based on the expression of personal opinions.
- The sheaf Laplacian $L_{\mathcal{F}} : C^0 \rightarrow C^0$ registers the “discord” in the system. The value of $L_{\mathcal{F}}(x)$ on each vertex v represents the difference between x_v (the opinions of v) and the opinions which would bring v in harmony with all its neighbors.
- The zeroth cohomology $H^0(G; \mathcal{F})$ computes the vector space of global sections — the space of opinion distributions in which all expressions of opinions are in harmony. As noted in Theorem 4.1, these are literally *harmonic* (in the kernel of the Laplacian).
- Relative cohomology $H^0(G, A; \mathcal{F})$ with respect to a subgraph $A \subset G$ represents harmonic opinion distributions which vanish on A . This is a measure of independence of the agents in A from the rest of the social network.

The discourse sheaf model is, in one sense, a mild generalization of the usual consensus problem over graphs. However, the ability to *program* a sheaf with linear transformations permits a number of features not present in the literature. Consider the simple example of three agents A, B, and C, all in pairwise communication.

- Because stalks vary from vertex-to-vertex, the discourse sheaf permits agents to have opinions on private basis topics. Agents A, B, and C need not have *any* basis topics in common.
- Because edge stalks are not identical to vertex stalks, the discourse sheaf model does not require everyone to share all their opinions with every neighbor; indeed, the topics for discussion need not relate at all to basis opinions of agents. Agents A and B might be discussing whether to eat lunch at the nearby pub. The edge stalk (favorability of dining at the pub) may not be something on which either A or B has a basis opinion.
- The restriction maps allow for the formation of *policies* from *principles*. For example, if agent A has a strong basis-opinion preference for sandwiches and is neutral about noise, the restriction map to the edge stalk could express a preference for dining at the (noisy, sandwich-renowned) pub. Agent B, who has basis opinions about walking long distances (dislikes) and quick meals (prefers) might have a restriction map that expresses dislike for the (not nearby) pub.
- Positive scalar multiplication acts both on vectors (intensifying or damping opinions) and on restriction maps (exaggerating or modulating expressed opinions). Negative scalar multiplication in a restriction map permits falsehoods: one can model agents who lie. Such dissembling or deception can be done selectively. What C says to B need not match what C says to A (even if they are discussing the same topic).
- There are multiple ways to set up opinion dynamics on a discourse sheaf. We begin, following the classical literature, by having individual agents change their opinions over time. This is perhaps not how real people engage in discourse. A different mode of evolution would permit *expression* of opinions to change, in order to bring discourse to a more harmonious state. This is achievable in the sheaf model by setting up dynamics on the restriction maps. Co-evolution of both opinion and expression is achievable in this model.

4. Sheaf Diffusion. Just as the graph Laplacian forms the basis for simple linear opinion dynamics, so does the sheaf Laplacian on discourse sheaves. Consider

the heat equation

$$(4.1) \quad \frac{dx}{dt} = -\alpha L_{\mathcal{F}}x, \quad : \quad \alpha > 0$$

on $x \in C^0(G; \mathcal{F})$. That is, x represents an opinion distribution, where $x_v \in \mathcal{F}(v)$ is the opinion of individual v . The diffusion dynamics tend to push the value at a node v toward greater agreement (as measured by the sheaf structure) with the expressed opinions of its neighbors. Our first result is that trajectories of this sheaf heat equation converge to global sections.

THEOREM 4.1. *Solutions $x(t)$ to (4.1) converge as $t \rightarrow \infty$ to the orthogonal projection of $x(0)$ onto $H^0(G; \mathcal{F})$.*

Proof. The sheaf Laplacian $L_{\mathcal{F}}$ is symmetric and positive semidefinite, and hence is diagonalizable with all eigenvalues nonnegative. The solution to (4.1) is

$$x(t) = \exp(-t\alpha L_{\mathcal{F}})x(0).$$

This solution operator has limit $\lim_{t \rightarrow \infty} \exp(-t\alpha L_{\mathcal{F}})$ equal to zero except on a block-diagonal identity submatrix corresponding to the zero eigenvalues. This is orthogonal projection onto $\ker L_{\mathcal{F}} = H^0(G; \mathcal{F})$. \square

Theorem 4.1 has several consequences. One is that the only stable opinion distributions are global sections of the discourse sheaf. If this sheaf has no nontrivial global sections, the only stable opinions will be everywhere zero: an uninteresting solution. Further, opinions converge exponentially to a stable distribution, with rate of convergence related to the spectral properties of the sheaf Laplacian [23].

The corresponding result for the graph Laplacian-based dynamics is that opinions converge toward the average of the initial opinion distribution. Discrete-time local averaging dynamics (with appropriately connected graphs) display a bit more flexibility, but admit only a limiting distribution that is a weighted average of the initial state. With sheaf Laplacian dynamics, new stable distributions are possible.

Example 4.2. (In polite company) The simple sheaf shown in Figure 4 has all stalks of dimension one and all restriction maps of full rank: each person [vertex] has a private opinion about a topic and expresses that opinion. For the sake of illustration, assume that each stalk (vertex and edge) has the same basis topic — say, opinion about a certain politician. If this were a constant sheaf, a global section would represent consensus with identical opinions. However, as illustrated, the two agents on the left have a positive personal opinion, whereas the two on the right have a negative personal opinion. The restriction maps encode expression of that opinion. Note how the discourse sheaf as illustrated permits selective expression of opinion. The two agents on the right tell a polite lie to their neighbors on the left but are frank with each other. A global section of this sheaf maintains the public agreement.

One notes that although each agent *knows* the veracity of their opinion expressions, they do not know the veracity of their neighbors' expressed opinions. Note also that, given this structure on the discourse sheaf, any initial opinion distribution will converge to one of these polarized distributions by Theorem 4.1. Compare the results of [2, 3], which studied similar structures implicitly.

5. Stubbornness and Harmonic Extension. Consider a slight variation on the diffusion equation (4.1) in which some agents are stubborn: they do not change their opinions in response to communication with their neighbors. What consequence




FIG. 4. A sheaf supporting stable polarized opinion distributions as global sections. The two agents on the right lie to their neighbors on the left, but are truthful to each other.

does this have for the long-run dynamics? To answer this question, we consider the problem of *harmonic extension* for partially-defined cochains on a sheaf. Our results here and in the next two sections are extensions of ideas originally introduced by Taylor [36] to the setting of discourse sheaves.

Let $U \subset V$ be a subset of vertices of G , and let $u \in C^0(U; \mathcal{F})$ be a 0-cochain with support in U — a choice of a data $u_v \in \mathcal{F}(v)$ for each $v \in U$. A harmonic extension of u to the rest of the graph is a 0-cochain $x \in C^0(G; \mathcal{F})$ such that $x|_U = u$ and $(L_{\mathcal{F}}x)_v = 0$ for every $v \in V \setminus U$. Harmonic extensions always exist; when $H^0(G, U; \mathcal{F}) = 0$, they are unique (see [23, Proposition 4.1] for a proof).

THEOREM 5.1. *The U -restricted dynamics*

$$(5.1) \quad \frac{dx}{dt} \Big|_v = \begin{cases} -\alpha(L_{\mathcal{F}}x)_v & : v \notin U \\ 0 & v \in U \end{cases}$$

on an initial condition x_0 converges exponentially to the harmonic extension of $u = (x_0)|_U$ nearest to x_0 .

Proof. As $u = x|_U$ is constant, we can rewrite the dynamics as acting purely on $y = x|_Y$ for $Y = V \setminus U$. Using $L_{\mathcal{F}}[\cdot, \cdot]$ to denote the block submatrix restricted to the indicated vertex sets, we can express the dynamics on $y(t)$ as

$$(5.2) \quad \frac{dy}{dt} = -\alpha(L_{\mathcal{F}}[Y, Y]y + L_{\mathcal{F}}[Y, U]u).$$

The fixed points of this dynamical system are the 0-cochains x where $L_{\mathcal{F}}x = 0$ on Y and $u = (x_0)|_U$ — the harmonic extensions of u . As $L_{\mathcal{F}}[Y, Y]$ is a principal submatrix of a positive semidefinite matrix, it is positive semidefinite; if $H^0(G, Y; \mathcal{F}) = 0$, it is positive definite. Further, $\text{im } L_{\mathcal{F}}[Y, U] \subseteq \text{im } L_{\mathcal{F}}[Y, Y] \perp \ker L_{\mathcal{F}}[Y, Y]$, so $\frac{dy}{dt}$ is always orthogonal to $\ker L_{\mathcal{F}}[Y, Y]$, and hence the dynamics preserve $\ker L_{\mathcal{F}}[Y, Y]$. Therefore, without loss of generality we can consider the system restricted to $\text{im } L_{\mathcal{F}}[Y, Y]$. That is, write $y = y^\perp + y^\parallel$, where $y^\parallel \in \text{im } L_{\mathcal{F}}[Y, Y]$ and $y^\perp \in \ker L_{\mathcal{F}}[Y, Y]$.

We now apply the general solution to an inhomogeneous linear ODE for x^\parallel to obtain

$$\begin{aligned} y(t) &= e^{-t\alpha L_{\mathcal{F}}[Y, Y]} y_0^\parallel - \int_0^t e^{-(t-\tau)\alpha L_{\mathcal{F}}[Y, Y]} \alpha L_{\mathcal{F}}[Y, U] u d\tau \\ &= e^{-t\alpha L_{\mathcal{F}}[Y, Y]} y_0^\parallel - \alpha^{-1} L_{\mathcal{F}}[Y, Y]^\dagger (I - e^{-\alpha L_{\mathcal{F}}[Y, Y]t}) \alpha L_{\mathcal{F}}[Y, U] u \\ &= e^{-t\alpha L_{\mathcal{F}}[Y, Y]} y_0^\parallel - L_{\mathcal{F}}[Y, Y]^\dagger (I - e^{-\alpha L_{\mathcal{F}}[Y, Y]t}) L_{\mathcal{F}}[Y, U] u. \end{aligned}$$

Here $L_{\mathcal{F}}[Y, Y]^{\dagger}$ is the Moore-Penrose pseudoinverse of $L_{\mathcal{F}}[Y, Y]$, which when restricted to $\text{im } L_{\mathcal{F}}[Y, Y]$ is simply the inverse. As $t \rightarrow \infty$, this expression converges to

$$(5.3) \quad y_{\infty}^{\parallel} = -L_{\mathcal{F}}[Y, Y]^{\dagger} L_{\mathcal{F}}[Y, U]u.$$

This is the minimum-norm harmonic extension of u , since any harmonic extension satisfies $L_{\mathcal{F}}[Y, Y]y + L_{\mathcal{F}}[Y, U]u = 0$. Since y_0^{\perp} is unchanged throughout, the limit is therefore $y_{\infty} = y_0^{\perp} + y_{\infty}^{\parallel}$, which clearly still satisfies the equation for harmonic extension of u . \square

This result quantifies how even a few stubborn individuals can have an influence on the opinion distribution throughout an entire social network. The opinion distribution is kept in constant tension induced by the stubborn individuals. We can characterize this equilibrium as the configuration minimizing total disagreement given the stubborn agents' opinions. That is, the harmonic extension is a minimizer of $\|\delta x\|^2$ subject to $x|_U = u$. If there is a global section x with $x|_U = u$, it is clearly a harmonic extension of u . However, harmonic extensions are not in general global sections.

6. Controlling Opinions. The analysis of stubborn individuals prompts the notion of more intentional control of opinions over a social network. By way of merging the notation of the previous section with that of linear controls, let $U \subset V$ be a set of *user input* vertices on which the controller can effect influence, and let $Y \subset V$ be a set of *observables* whose preferences one wants to control. We will denote by $u(t) \in C^0(U; \mathcal{F})$ and $y(t) \in C^0(Y; \mathcal{F})$ data supported on U and Y respectively taking values in the stalks of \mathcal{F} on those respective vertex sets.

The influence of user inputs on the system is mediated through a linear transformation $B : C^0(G; \mathcal{F}) \rightarrow C^0(G; \mathcal{F})$ with image and coimage $C^0(U; \mathcal{F})$; the observables are viewed through $C : C^0(G; \mathcal{F}) \rightarrow C^0(G; \mathcal{F})$ with image and coimage $C^0(Y; \mathcal{F})$. The resulting linear control system is:

$$(6.1) \quad \frac{dx}{dt} = -\alpha L_{\mathcal{F}}x + Bu \quad : \quad \frac{dy}{dt} = Cx.$$

Controllability of (6.1) answers the natural question of whether opinion distributions on Y can be determined via manipulation of inputs on U . The system will naturally settle on a stable opinion distribution — a global section of \mathcal{F} . Do inputs exist that will steer the system to an arbitrary global section?

Consider the simplified case where B is the identity map on $C^0(U; \mathcal{F})$ and zero elsewhere. In this case, we have the following result:

THEOREM 6.1. *If the relative cohomology $H^0(G, U; \mathcal{F}) = 0$, then the system (6.1) is stabilizable, with B the identity on $C^0(U; \mathcal{F})$.*

Proof. Stabilizability is equivalent to the condition that the matrix

$$\begin{bmatrix} (-\alpha L_{\mathcal{F}} - \lambda I) & B \end{bmatrix}$$

be full rank for all λ with nonnegative real part (see, e.g., [33] for a deeper discussion of this and other standard results of linear control theory). Since the eigenvalues of $-\alpha L_{\mathcal{F}}$ are real and nonpositive, $-\alpha L_{\mathcal{F}} - \lambda I$ is already full rank for any λ with nonzero imaginary part or negative real part. Thus we only need to consider $\lambda = 0$, the case of the matrix $\begin{bmatrix} -\alpha L_{\mathcal{F}} & B \end{bmatrix}$. This matrix has full rank if for every $x \in \ker L_{\mathcal{F}}$, there exists some $u \in \text{im } B$ with $x^T u \neq 0$. In particular, this will be satisfied if no nontrivial global section of \mathcal{F} vanishes on U . From the definition and interpretation of relative cohomology, this is precisely the condition that $H^0(G, U; \mathcal{F}) = 0$. \square




FIG. 5. *Weighted reluctance to changing opinions may be modeled by extending the discourse sheaf over an augmented graph, giving each vertex a stubborn parent who exerts influence.*

This theorem implies that given an appropriate input set, we can ensure that the dynamics converge to any global section of \mathcal{F} . When $H^0(G, U; \mathcal{F}) = 0$, every global section is the unique harmonic extension of its restriction to U , so we need only control $u = x|_U$ to the desired states, and the rest of the network will follow.

The dual result to Theorem 6.1 (presented without proof) is

THEOREM 6.2. *If the relative cohomology $H^0(G, Y; \mathcal{F}) = 0$, then the system (6.1) is detectable for C the identity on $C^0(Y; \mathcal{F})$ and zero elsewhere.*

Thus, by observing agents on Y we can deduce any motion of the global state projected to $H^0(G; \mathcal{F})$. Given both conditions we can construct an observer-controller pair steering the system to a global section with any desired outcome as measured on Y .

Example 6.3. If G is connected and the communication structure is given by the constant sheaf \mathbb{R}^n , any vertex gives an input set for which the dynamics are stabilizable. This is because there are no nonzero constant \mathbb{R}^n -valued functions on G that vanish at a vertex, so $H^0(G, \{v\}; \mathbb{R}^n) = 0$. In terms of opinion dynamics, it is only necessary to have arbitrary influence on a single individual in order to ensure eventual global consensus on any given opinion — a trivial property of the constant sheaf.

7. Weighted Reluctance. The control perspective is helpful in analyzing another type of linear dynamics, where individuals are resistant to modifying their initial opinions (but not infinitely so, as in Section 5). We model this as a feedback controller attached to the original system, letting $u_v = \alpha\gamma_v((x_0)_v - x_v)$, where γ_v is an agent-dependent reluctance parameter. The long-run opinion distribution is again a function of harmonic extension. This can be effected by expanding the sheaf to an augmented graph G' as follows (see Figure 5).

Construction: Given G , augment it to a graph G' by duplicating the vertex set V to a copy V' . Attach each $v' \in V'$ to the corresponding vertex $v \in G$, with a single edge e' . Extend the discourse sheaf \mathcal{F} to a sheaf \mathcal{F}' on G' by letting $\mathcal{F}'(v') = \mathcal{F}'(e') = \mathcal{F}(v)$, where e' is the edge between v' and v . The restriction maps are $\mathcal{F}'_{v \triangleleft e'} = \mathcal{F}'_{v' \triangleleft e'} = \sqrt{\gamma_v}I$. One thinks of this augmented graph as giving each agent an additional acquaintance, their *parent*, who acts as a constant influence on their opinions.

We now apply the results about dynamics with stubborn agents to \mathcal{F}' on G' , with parents as stubborn agents. For an initial condition $x_0 \in C^0(G; \mathcal{F})$, extend it to G' via $x_0(v') = x_0(v)$ and force all the agents v' to be stubborn (in the sense of §5). The

differential equation governing the opinion evolution is

$$\begin{aligned}
 (7.1) \quad \frac{dx_v}{dt} &= -\alpha(L_{\mathcal{F}'}x)_v \\
 &= -\alpha((L_{\mathcal{F}}x)_v + \gamma_v(x_v - x_{v'})) \\
 &= -\alpha(L_{\mathcal{F}}x)_v + \alpha\gamma_v((x_0)_v - x_v).
 \end{aligned}$$

Applying Theorem 5.1, we see that the opinion distribution converges to the harmonic extension of $(x_0)|_{V'}$ to the rest of \mathcal{F}' . The stubbornness parameters γ_v influence how much the initial opinions influence the limit, and hence how far from a global section of \mathcal{F} the limiting opinion distribution lies.

8. Learning to Lie. The notion that communication over a social network leads to changes in opinions is a convenient idealization to which diffusion dynamics applies. Because using a discourse sheaf to model opinion dynamics makes explicit the communication structure employed by the agents, it allow us to model changes in that structure. Instead of opinions changing over time, one could just as well consider evolution of expression: agents can learn to communicate differently based on the reactions of their neighbors. Leaving aside the sociological questions of whether a typical person in the face of opposition actually changes opinions or simply “learns to communicate better,” we demonstrate the flexibility of the discourse sheaf model under such settings.

Assume that each agent v is able to modify all its restriction maps $\mathcal{F}_{v \triangleleft e}$ (for edges e incident to v) and is able to observe their neighbors’ translated opinions $\mathcal{F}_{u \triangleleft e}x_u$. If the goal of each agent is to learn how to translate their opinions so that apparent consensus is reached, they should alter their restriction maps to remove the part of the image that contributes to disagreement with neighbors. That is, the dynamics should be of the form

$$(8.1) \quad \frac{d}{dt}\mathcal{F}_{v \triangleleft e} = -\beta(\mathcal{F}_{v \triangleleft e}x_v - \mathcal{F}_{u \triangleleft e}x_u)x_v^T,$$

for some diffusion strength $\beta > 0$. These dynamics may be nicely expressed in terms of the block rows δ_e of the coboundary matrix corresponding to each edge:

$$(8.2) \quad \frac{d\delta_e}{dt} = -\beta\delta_e x_e x_e^T,$$

where x_e is the vector in $C^0(G; \mathcal{F})$ in which all entries corresponding to vertices not incident to e have been replaced with zero. Combining these together, we have

$$(8.3) \quad \frac{d\delta}{dt} = -\beta P_\delta(\delta x x^T),$$

where P_δ is the map that takes the matrix of a linear transformation $C^0(G; \mathcal{F}) \rightarrow C^1(G; \mathcal{F})$ and projects it to a matrix with the correct sparsity pattern to be a sheaf coboundary matrix. That is, P_δ sets all entries for blocks corresponding to non-incident vertex-edge pairs to zero.

THEOREM 8.1. *For $\mathcal{F}(t)$ a solution to (8.1) on the space of sheaves over G with fixed stalks, the sheaf $\mathcal{F} = \mathcal{F}(0)$ converges to $\mathcal{F}' = \lim_{t \rightarrow \infty} \mathcal{F}(t)$, the nearest sheaf such that x is a global section, where distance between \mathcal{F} and \mathcal{F}' is measured by the squared Frobenius norm:*

$$(8.4) \quad d(\mathcal{F}, \mathcal{F}') = \sum_{v \triangleleft e} \|\mathcal{F}_{v \triangleleft e} - \mathcal{F}'_{v \triangleleft e}\|_F^2 = \|\delta - \delta'\|_F^2.$$




FIG. 6. A constant sheaf over an edge with an initial opinion distribution that is highly discordant [left] converges under diffusion of the sheaf to a nonconstant sheaf [right] in which the agent with the more extreme negative opinion has “learned to lie” in order to come to consensus.

Proof. The 0-cochain x is a global section of \mathcal{F}' precisely when, for each edge $e = u \sim v$, $\mathcal{F}'_{u \triangleleft e} x_u = \mathcal{F}'_{v \triangleleft e} x_v$, or equivalently, when $\delta'_e x_e = 0$.

The dynamics for each edge are uncoupled, as can be seen from the form of the equation in (8.2). The differential equation for each edge is linear, given by the operator A taking δ_e to $\delta_e x_e x_e^T$. This operator is positive semidefinite with respect to the inner product $\langle \delta_e, \delta'_e \rangle = \text{tr}(\delta_e^T \delta'_e)$, and its kernel is given by those δ_e for which $\delta_e x_e = 0$. Therefore, the same argument as in the proof of Theorem 4.1 shows that the trajectory of δ_e converges to its orthogonal projection onto $\ker A$ with respect to this inner product. The corresponding norm is the Frobenius norm, and hence δ_e converges to the nearest δ'_e as measured by this norm such that $\delta'_e x_e = 0$.

Since the distance (8.4) decomposes across edges, and each edge independently satisfies this distance-minimizing property, the same is true of their combination into a complete coboundary matrix. That is, the limiting coboundary matrix $\delta' = \lim_{t \rightarrow \infty} \delta(t)$ is the minimizer of $\|\delta(0) - \delta'\|_F^2$ such that δ' is a sheaf coboundary matrix and $\delta' x = 0$. \square

This proof indicates a sort of duality between the sheaf heat equation (4.1) and the structural dynamics (8.1). Both are diffusion-like processes, adjusting parameters to alleviate a local discrepancy.

Example 8.2. (Learning to lie) The restriction map diffusion dynamics can convert the constant sheaf into a nontrivial communication structure. Start with the constant sheaf on a graph with two vertices and a single edge as shown in Figure 6, with a highly inconsistent 0-cochain assigning -4 to one vertex and $+1$ to the other. As the restriction map dynamics progress, the discourse sheaf becomes an inconsistent one, and the sign of one restriction map changes — the corresponding agent *learns to lie* about their opinion. Note that this agent changes their expressed opinion’s sign, but also downplays its magnitude, allowing the original opinion distribution to become a global section.

9. Joint Opinion-Expression Diffusion. The natural culmination of this line of reasoning is to combine the opinion diffusion (4.1) and the communication structure diffusion (8.1). That is, evolve both the 0-cochain x and the restriction maps $\mathcal{F}'_{v \triangleleft e}$ according to

$$(9.1) \quad \begin{aligned} \frac{dx}{dt} &= -\alpha \delta^T \delta x \\ \frac{d\delta_e}{dt} &= -\beta \delta_e x_e x_e^T. \end{aligned}$$

LEMMA 9.1. *For x, δ evolving according to (9.1), the function*

$$(9.2) \quad \Psi(x, \delta) = \frac{1}{2} x^T \delta^T \delta x$$

satisfies $\Psi(x, \delta) \geq 0$ and $\frac{d}{dt} \Psi(x, \delta) \leq 0$, with zero attained in both instances if and only if $\delta x = 0$.

Proof. In (9.1), δ is considered to lie in the space of linear transformations with the appropriate sparsity pattern to be the coboundary of a sheaf over G . That $\Psi(x, \delta) \geq 0$ is immediate from its definition, as is its vanishing precisely when $\delta x = 0$. The proof of the second inequality is by direct computation:

$$(9.3) \quad \begin{aligned} \frac{d}{dt} \Psi(x, \delta) &= x^T \delta^T \delta \frac{dx}{dt} + x^T \delta^T \frac{d\delta}{dt} x \\ &= -\alpha x^T \delta^T \delta \delta^T \delta x - \beta x^T \delta^T P_\delta (\delta x x^T) x \\ &= -\alpha x^T (\delta^T \delta)^2 x - \beta \sum_e x_e^T \delta_e^T \delta_e x_e x_e^T x_e. \end{aligned}$$

The first term is clearly nonpositive, and negative precisely when $\delta x \neq 0$, and the second term is similarly negative when $\delta_e x_e \neq 0$ for some e . \square

Applying LaSalle's invariance principle to this function reveals that all limit points of trajectories of (9.1) satisfy $\delta x = 0$. More is true: it is easy to check that (9.1) is actually a gradient descent equation on Ψ , with the gradient defined with respect to the inner product $\langle (x, \delta), (x', \delta') \rangle = \frac{1}{\alpha} x^T x' + \frac{1}{\beta} \sum_e \text{tr}(\delta_e^T \delta'_e)$. This implies that trajectories indeed converge to points with $\delta x = 0$. That is, trajectories converge to points describing a sheaf \mathcal{F} on G with coboundary map δ together with a global section x of \mathcal{F} .

Because the stationary points of (9.1) are not isolated, there is no global asymptotic stability. The set of equilibria is the set of solutions to a system of degree-2 polynomials forming a singular algebraic variety. The equilibria lying at nonsingular points of this variety are, however, Lyapunov stable.

THEOREM 9.2. *The smooth stationary points of (9.1) are Lyapunov stable.*

Proof. It remains to show that given a stationary point (x_*, δ_*) at which the derivative of the map $(x, \delta) \mapsto \delta x$ is full rank, any neighborhood of (x_*, δ_*) contains the forward orbit of a subneighborhood. In a neighborhood of (x_*, δ_*) , the set of equilibria of (9.1) is a smooth manifold, given by the equation $\delta x = 0$. The tangent space at this point is the kernel of the map $(\delta, x) \mapsto \delta_* x + \delta x_*$. Meanwhile, the linearization of (9.1) about (x_*, δ_*) is

$$\begin{aligned} \frac{dx}{dt} &= -\alpha (\delta_*^T \delta x_* + \delta_*^T \delta_* x) \\ \frac{d\delta_e}{dt} &= -\beta [\delta_e^T (x_*)_e (x_*)_e^T + (\delta_*)_e^T x_e (x_*)_e^T] \end{aligned}$$

These are zero if and only if (x, δ) satisfy $\delta_* x + \delta x_* = 0$, so the stationary set of the linearization is precisely the same as the tangent space at (x_*, δ_*) . Therefore, the stationary points near (x_*, δ_*) describe a center manifold for the dynamics. Standard results on stability of the center manifold indicate that the stability of an equilibrium within the center-stable manifold is equivalent to its stability in the center manifold [28]. Since (x_*, δ_*) is Lyapunov stable within the center manifold (because the dynamics are trivial), it is therefore Lyapunov stable. \square

The regions of highest dimension of the variety X are those where $\ker \delta = 0$ (for most choices of G and stalk dimensions). This would seem to indicate that we should expect most trajectories to converge to points with $x = 0$. However, this is not the case.

THEOREM 9.3. *If one of the diagonal blocks of $\alpha\delta_0^T\delta_0 - \beta x_0^T x_0$ fails to be semi-definite, the trajectory of (9.1) converges to a point $(x_\infty, \delta_\infty)$ with $x_\infty \neq 0$.*

Proof. Let $M = \delta^T \delta - x x^T$ and consider

$$\begin{aligned}
(9.4) \quad \frac{dM}{dt} &= \frac{d}{dt} \alpha \delta^T \delta - \beta x x^T \\
&= \alpha [-\beta P_\delta (\delta x x^T)^T \delta - \beta \delta^T P_\delta (\delta x x^T)] - \beta (-\alpha x x^T \delta^T \delta - \alpha \delta^T \delta x x^T) \\
&= -\alpha \beta [(\delta^T P_\delta (x x^T \delta x x^T))^T + \delta^T P_\delta (\delta x x^T)] + \alpha \beta [(\delta^T \delta x x^T)^T + \delta^T \delta x x^T].
\end{aligned}$$

Note that the diagonal block of $\delta^T P_\delta (\delta x x^T)$ corresponding to a vertex v is equal to $(\delta^T \delta x x^T)_{v,v}$, since the sparsity pattern of the relevant block row of δ^T is the same as the sparsity pattern imposed by the projection P . Thus, when we restrict to the diagonal blocks, the derivatives cancel and we have $\frac{d}{dt} \text{diag}(M) = 0$. Therefore, if $M_{v,v}$ is indefinite at $t = 0$, it must be indefinite for all t . In particular, this means that $x x^T$ cannot approach zero, since otherwise M and all of its diagonal subblocks would approach semidefiniteness. \square

This condition implies that if any diagonal element of M is negative, the dynamics converge to a nonzero x . *A fortiori*, if $\text{tr}(M) = \|\delta\|_F^2 - \|x\|^2$ is negative, the limiting value of x is nonzero. Thus, given any initial δ_0 and x_0 , there exists some scaling factor κ such that the initial condition $(\delta_0, \kappa x_0)$ converges to a sheaf with a nontrivial global section.

Other relevant quantities are controlled during the combined diffusion dynamics.

THEOREM 9.4. *The quantities $\|\delta\|_F^2$, $\|x\|^2$, $\|\delta x\|^2$, and $\frac{\|\delta x\|^2}{\|x\|^2}$ are nonincreasing under the evolution of (9.1).*

Proof. That $\|\delta x\|^2$ is decreasing is implied by the fact that Ψ is decreasing under (9.1). We evaluate the other derivatives:

$$\begin{aligned}
(9.5) \quad \frac{d}{dt} \|\delta\|_F^2 &= \frac{d}{dt} \text{tr}(\delta^T \delta) \\
&= -\beta \text{tr}(\delta^T P_\delta (\delta x x^T) + P_\delta (\delta x x^T)^T \delta) \\
&= -2\beta \sum_e (x_e^T \delta_e^T \delta_e x_e) x_e^T x_e \leq 0,
\end{aligned}$$

$$\begin{aligned}
(9.6) \quad \frac{d}{dt} \|x\|^2 &= \frac{d}{dt} x^T x \\
&= -\alpha (x^T \delta^T \delta x + (\delta^T \delta x)^T x) \\
&= -2\alpha x^T \delta^T \delta x \leq 0.
\end{aligned}$$

Finally,

$$(9.7) \quad \frac{d}{dt} \frac{\|\delta x\|^2}{\|x\|^2} = \frac{\|x\|^2 \frac{d}{dt} \|\delta x\|^2 - \|\delta x\|^2 \frac{d}{dt} \|x\|^2}{\|x\|^4}.$$




FIG. 7. A constant sheaf over an edge with an initial opinion distribution that is highly discordant [left] converges under the combined opinion-expression diffusion to a nonconstant sheaf [right] in which both opinions and expressions have relaxed to come to consensus. Under this initial condition, the agent on the left has learned to lie.

We know that $\frac{d}{dt} \|\delta x\|^2 \leq -\alpha x^T (\delta^T \delta)^2 x$, so this is bounded above by

$$\frac{-2\alpha \|x\|^2 (\alpha x^T (\delta^T \delta)^2 x) + 2\alpha (x^T \delta^T \delta x)^2}{\|x\|^4} = 2\alpha \left[\left(\frac{x^T \delta^T \delta x}{\|x\|^2} \right)^2 - \frac{x^T (\delta^T \delta)^2 x}{\|x\|^2} \right].$$

Thus, $\frac{\|\delta x\|^2}{\|x\|^2}$ is decreasing if the inequality

$$\left(\frac{x^T \delta^T \delta x}{\|x\|^2} \right)^2 \leq \frac{x^T (\delta^T \delta)^2 x}{\|x\|^2}$$

holds. By taking $\|x\| = 1$ and diagonalizing $\delta^T \delta$, we get the equivalent inequality

$$\sum_i \lambda_i^2 x_i \geq \left(\sum_i \lambda_i x_i \right)^2$$

for $\lambda_i, x_i \geq 0$, $\sum_i x_i = 1$, which is simply Jensen's inequality. \square

This last decreasing observable is the Rayleigh quotient of $L = \delta^T \delta$ corresponding to x . That it is decreasing means $x(t)$ becomes a global section of $\mathcal{F}(t)$ at least as quickly as it approaches zero.

Example 9.5. (Learning to lie, redux) The combined diffusion dynamics also enable an agent to learn to falsify opinions. Consider the same initial conditions as before, but run the combined dynamics with $\alpha = \beta = 1$. The sheaf and cochain converge again to a sheaf where one agent lies about their opinion. The opinion distribution changes somewhat, but because the discrepancy between opinions is so great, the sheaf is able to adapt more easily than the opinions themselves: see Figure 7.

10. Nonlinear Laplacians. The sheaf Laplacian is constructed from the sheaf coboundary and its adjoint, with an implicit isomorphism between C^1 and its dual given by the standard inner product on the edge stalks. We can make this mapping explicit, and insert a new function between the terms in order to produce nonlinear Laplacians with new behaviors.

Let $\phi_e : \mathcal{F}(e) \rightarrow \mathcal{F}(e)$ be a continuous but not-necessarily-linear map for each edge e of G , and define $\Phi : C^1(G; \mathcal{F}) \rightarrow C^1(G; \mathcal{F})$ by combining these edgewise maps. The corresponding nonlinear Laplacian is $L_{\mathcal{F}}^{\Phi} = \delta_{\mathcal{F}}^T \circ \Phi \circ \delta_{\mathcal{F}}$. Because the nonlinear map Φ is applied edgewise, $L_{\mathcal{F}}^{\Phi} x$ can still be computed locally in the network. Thus, the nonlinear heat equation

$$(10.1) \quad \frac{dx}{dt} = -\alpha L_{\mathcal{F}}^{\Phi} x$$

over a discourse sheaf \mathcal{F} describes a form of networked opinion dynamics.

One way to construct a nonlinear Laplacian is by beginning with a set of edge potential functions. The heat equation on a sheaf is gradient descent with respect to x on the potential function $\Psi(x) = \frac{1}{2}\|\delta x\|^2 = \sum_e \frac{1}{2}\|\delta_e x\|^2$. We can replace this potential function with a new function defined edgewise:

$$\Psi(x) = \sum_e U_e(\delta_e x) = U(\delta x),$$

for once-differentiable edgewise potential functions $U_e : \mathcal{F}(e) \rightarrow \mathbb{R}$. The gradient of this function at a point x is $\nabla\Psi(x) = \delta^T \circ \nabla U \circ \delta x$. This is therefore a nonlinear sheaf Laplacian $L_{\mathcal{F}}^{\Phi}$ with $\Phi = \nabla U$. The heat equation for this nonlinear Laplacian is precisely gradient descent on Ψ .

Analysis of the heat equation is simplified for nonlinear Laplacians constructed from edge potentials. For instance, if the potential functions U_e are convex, Ψ serves as a Lyapunov function, ensuring stability of the dynamics. If each U_e has a local minimum at 0, global sections of \mathcal{F} will be stationary points of the heat equation. Indeed, if each U_e is radially unbounded with a single local minimum at 0, the long-term behaviors of the linear and nonlinear heat equations agree.

PROPOSITION 10.1. *For each edge e of G , let $U_e : \mathcal{F}(e) \rightarrow \mathbb{R}$ be a differentiable, radially unbounded function with a unique local minimum at 0. Trajectories of the nonlinear heat equation*

$$\frac{dx}{dt} = -\alpha L_{\mathcal{F}}^{\nabla U} x$$

converge to the orthogonal projection of the initial condition onto $H^0(G; \mathcal{F})$.

Proof. First observe that $\frac{dx}{dt} \in \text{im } \delta^T$ and hence is always orthogonal to $\ker \delta = H^0(G; \mathcal{F})$. Decomposing $x = x^{\parallel} + x^{\perp}$, where x^{\parallel} is the orthogonal projection of x onto $H^0(G; \mathcal{F})$, we restrict our attention to $\text{im } \delta^T$ and the evolution of x^{\perp} .

Letting $\Psi(x^{\perp}) = U(\delta x^{\perp})$, we have a function vanishing precisely when $\delta x^{\perp} = 0$, which holds precisely when $x^{\perp} = 0$. Further,

$$\frac{d\Psi}{dt} = \langle \nabla\Psi(x^{\perp}), \frac{dx^{\perp}}{dt} \rangle = \langle L_{\mathcal{F}}^{\nabla U} x^{\perp}, -\alpha L_{\mathcal{F}}^{\nabla U} x^{\perp} \rangle \leq 0,$$

with equality precisely when $L_{\mathcal{F}}^{\nabla U} x^{\perp} = 0$. Because U is radially unbounded and has a unique local minimum at 0, ∇U vanishes only at 0, and hence $L_{\mathcal{F}}^{\nabla U} x^{\perp} = 0$ if and only if $x^{\perp} = 0$. Thus $\frac{d\Psi}{dt}$ vanishes only at the origin and is therefore a Lyapunov function for the nonlinear heat equation restricted to $\text{im } \delta^T$. Radial unboundedness of U implies radial unboundedness of Ψ and therefore the origin is globally asymptotically stable, meaning $x^{\perp} \rightarrow 0$. Therefore, $\lim_{t \rightarrow \infty} x(t) = x^{\parallel}(0)$. \square

11. Bounded Confidence. Nonlinear Laplacian dynamics (and in particular edge potential dynamics) can be used to extend the popular *bounded confidence* opinion dynamics to discourse sheaves. The central idea behind bounded confidence opinion dynamics is that individuals only have confidence in the opinions of neighbors that are sufficiently similar to their own, and thus only take these opinions into account when updating. The reigning discrete-time model of bounded confidence dynamics is based on that of Hegselmann and Krause [25], with extensions to multidimensional opinions. In this model, each agent has a threshold D , and only pays heed to opinions of neighbors that are within distance D of their own opinion. Continuous-time

versions of this model have been analyzed, both with sharp discontinuous thresholds [6, 9] as well as smooth transitions between confidence and no-confidence [10].

We will here discuss a continuous-time version of bounded confidence dynamics for sheaves, using smooth threshold functions. These will be represented in terms of edgewise potential functions. Given a discourse sheaf \mathcal{F} on a graph G , for each edge e of G choose a threshold D_e and a differentiable function $\psi_e : [0, \infty) \rightarrow \mathbb{R}$ such that $\psi'_e(y) = 0$ for $y \geq D_e$ and $\psi'_e(y) > 0$ for $y < D_e$. The edge potential function $U_e : \mathcal{F}(e) \rightarrow \mathbb{R}$ is then given by $U_e(y_e) = \psi_e(\|y_e\|^2)$. The gradient of this potential is $\nabla U_e(y_e) = \psi'_e(\|y_e\|^2)y_e$, and therefore the associated nonlinear Laplacian is given by

$$(11.1) \quad L_{\mathcal{F}}^{\nabla U} x = \delta^T \text{diag}(\psi'_e(\|\delta_e x\|^2)) \delta x.$$

This formula can be written vertexwise as

$$(11.2) \quad (L_{\mathcal{F}}^{\nabla U} x)_v = \sum_{u, v \triangleleft e} \mathcal{F}_{v \triangleleft e}^T \psi'_e(\|\mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u\|^2) (\mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u).$$

In comparison with the formula for the standard sheaf Laplacian, there is a nonlinear scaling factor depending on the discrepancy over each edge. This nonlinear Laplacian $L_{\mathcal{F}}^{\nabla U}$ can be used to generate bounded confidence dynamics.

THEOREM 11.1. *Suppose that $U_e(y_e) = \psi_e(\|y_e\|^2)$ for some $\psi_e : [0, \infty) \rightarrow \mathbb{R}$ with $\psi'_e(y) = 0$ for $y \geq D_e$ and $\psi'_e(y) > 0$ for $y < D_e$. Then an opinion distribution $x \in C^0(G; \mathcal{F})$ is harmonic with respect to $L_{\mathcal{F}}^{\nabla U}$ if and only if for every edge $e = u \sim v$ with $\mathcal{F}_{v \triangleleft e} x_v \neq \mathcal{F}_{u \triangleleft e} x_u$, $\|\mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u\|^2 \geq D_e$.*

Proof. If $x \in H^0(G; \mathcal{F})$, then $\delta x = 0$, and since $\nabla U_e(0) = 0$, $\delta^T \nabla U(\delta x) = \delta^T 0 = 0$. More generally, $\nabla U_e(\delta x)_e = 0$ whenever either $(\delta x)_e = 0$ or $\|(\delta x)_e\|^2 \geq D_e$. But $(\delta x)_e = \mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u$.

Conversely, if $L_{\mathcal{F}}^{\nabla U}(x) = 0$, then $\nabla U(\delta x) \in \ker \delta^T$; equivalently, $\nabla U(\delta x)$ is orthogonal to $\text{im } \delta$. In particular, $\langle \nabla U(\delta x), \delta x \rangle$ must be zero. Letting $y = \delta x$, we have

$$(11.3) \quad \langle \nabla U(y), y \rangle = \sum_e \langle \nabla U_e(y_e), y_e \rangle = \sum_e \langle \psi'_e(\|y_e\|^2) y_e, y_e \rangle.$$

These terms are all nonnegative, and $\langle \psi'_e(\|y_e\|^2) y_e, y_e \rangle = 0$ if and only if either $y_e = 0$ or $\psi'_e(\|y_e\|^2) = 0$. The first condition holds precisely when $(\delta x)_e = 0$, and the second holds precisely when $\|(\delta x)_e\|^2 \geq D_e$. \square

Naturally, one constructs the bounded confidence Laplacian to study the corresponding diffusion dynamics

$$(11.4) \quad \frac{dx}{dt} = -L_{\mathcal{F}}^{\nabla U}(x).$$

Theorem 11.1 identifies the equilibria of these dynamics. Global sections of \mathcal{F} are still equilibria, but there are more. Given $x \in C^0(G; \mathcal{F})$, we construct a subgraph G_x of G and a sheaf \mathcal{F}_x on G_x as follows: G_x has the same vertices as G , but only contains the edges e where $\|(\delta x)_e\| < D_e$. The sheaf \mathcal{F}_x is the same as \mathcal{F} but with the data for edges not in G_x removed. So a 0-cochain $x \in C^0(G; \mathcal{F})$ is a fixed point for (11.4) if and only if it is a global section of \mathcal{F}_x . We denote the subset of $C^0(G; \mathcal{F})$ for which the corresponding sheaf is \mathcal{F}_x by K_x . That is,

$$(11.5) \quad K_x = \{x \in C^0(G; \mathcal{F}) : \|\delta_e x\|^2 < D_e \text{ for } e \in G_x, \|\delta_e x\|^2 \geq D_e \text{ for } e \notin G_x\}.$$

Suppose that x_0 is a fixed point with $\|(\delta x_0)_e\|^2 \neq D_e$ for all edges e . Sufficiently close to x_0 , the dynamics behave like the standard diffusion dynamics on \mathcal{F}_{x_0} .

THEOREM 11.2. *Let x_0 be a fixed point of (11.4) lying in the interior of K_{x_0} — that is, with $\|(\delta x_0)_e\| > D_e$ for all $e \notin G_{x_0}$. There exists a neighborhood W of x_0 such that for every initial condition $x \in W$, the trajectory of (11.4) converges to the nearest global section of \mathcal{F}_{x_0} .*

Proof. Let W be a neighborhood of x_0 contained in K_{x_0} satisfying:

1. If $x \in W$, its orthogonal projection x^\parallel onto $H^0(G_{x_0}; \mathcal{F}_{x_0})$ is in W .
2. If $x \in W$ is not a fixed point with $x = x^\parallel + x^\perp$, then $\|x^\perp\| < \frac{\|\delta_e x^\parallel\|^2 - D_e}{2\|\delta_e x^\parallel\|\|\delta_e\|}$ for all $e \notin G_{x_0}$.
3. If $x = x^\parallel + x^\perp \in W$ is not a fixed point, then $\|x^\perp\|^2 < \frac{D_e}{\|\delta_e\|^2}$ for all $e \in G_{x_0}$.

Such a neighborhood exists by continuity of the linear maps δ_e and because we can choose arbitrarily small tubular neighborhoods around the set of fixed points.

Consider $x = x^\parallel + x^\perp \in W$. We will show that x converges to x^\parallel , which is in U by condition (1). Note that as long as $x \in K_{x_0}$, $\frac{d}{dt}x^\parallel = 0$. Similarly, for $x \in K_{x_0}$, $\frac{d}{dt}\frac{1}{2}\|x^\perp\|^2 = \langle x^\perp, \frac{d}{dt}x^\perp \rangle = \langle x^\perp, -\delta^T \nabla U(\delta x^\perp) \rangle = -\langle \delta x^\perp, \nabla U(\delta x^\perp) \rangle \leq 0$, by the argument in the proof of Theorem 11.1. This is zero only if $x^\perp = 0$, and hence $\|x^\perp\|^2$ is strictly decreasing as long as it is nonzero and $x \in K_{x_0}$. In particular, for any initial condition $x \in W$ with $x^\perp \neq 0$ there is some maximal time interval $[0, T)$ on which $\frac{d}{dt}\|x^\perp\| < 0$ is strictly decreasing.

Combining this with (2) and (3) above reveals that on this time interval, for every $e \notin G_{x_0}$,

$$\begin{aligned} \|\delta_e x\|^2 &\geq \|\delta_e x^\parallel\|^2 - 2\|\delta_e x^\perp\|\|\delta_e x^\parallel\| \\ &\geq \|\delta_e x^\parallel\|^2 - 2\|\delta_e\|\|x^\perp\|\|\delta_e x^\parallel\| \\ &> \|(\delta x^\parallel)_e\|^2 - 2\|\delta_e\|\|\delta_e x^\parallel\| \frac{\|\delta_e x^\parallel\|^2 - D_e}{2\|\delta_e x^\parallel\|\|\delta_e\|} = D_e \end{aligned}$$

by condition (2). Similarly, for $e \in G_{x_0}$,

$$\|(\delta x)_e\|^2 = \|(\delta x^\perp)_e\|^2 \leq \|\delta_e\|^2 \|x^\perp\|^2 < D_e$$

by condition (3). Since these relations hold at time $t = 0$ and $\|x^\perp\|$ is decreasing, they hold for all $t \in [0, T)$ as well. This ensures that x remains in K_{x_0} and does not approach the boundary on this interval. If T is finite, we may thus conclude that $\frac{d}{dt}\|x^\perp\| < 0$ at $t = T$ as well, so it must be that $T = \infty$. Thus x remains in K_{x_0} and $\|x^\perp\|$ is strictly decreasing for all time.

It remains to show that trajectories converge: that $x^\perp \rightarrow 0$. This happens because $\Psi_{x_0}(x^\perp) = \sum_{e \in G_{x_0}} U_e(\delta_e x^\perp)$ is strictly decreasing along trajectories and vanishes precisely when $x^\perp = 0$. \square

12. Antagonistic Dynamics. One may wish to model certain agent pairs who, rather than attempting to move toward mutual consensus, instead try to increase the distance between their expressed opinions. Such a dynamical structure might correspond to relationships of distrust or enmity. We will call such a relationship a *negative edge* and partition the edge set of the graph G into negative edges, E_- , and complementary positive edges E_+ .

One way to model relationships over negative edges in a discourse sheaf is to change the sign of one restriction map on each negative edge. Instead of a tactful

lie, this might be interpreted as an “agreement to disagree” between the neighbors, making a stable pair of opinions one where the expressed opinions are the same in magnitude but opposite in sign. In the case where the discourse sheaf is simply a constant sheaf, this is essentially the approach taken in [2, 3].

A better approach is to leave the discourse sheaf untouched and modify the dynamics appropriately. Let E_- be the set of negative links between agents, and E_+ its complement, the set of positive links. We define edge potential functions associated to this signing:

$$U_e(y) = \begin{cases} \|y\|^2 & e \in E_+ \\ -\|y\|^2 & e \in E_- \end{cases}.$$

The corresponding nonlinear Laplacian $L_{\mathcal{F}}^{\nabla U}$ describes dynamics where agents attempt to move toward consensus as quickly as possible over positive edges and away from consensus as quickly as possible over negative edges. Let $S : C^1(G; \mathcal{F}) \rightarrow C^1(G; \mathcal{F})$ be the block diagonal matrix whose blocks corresponding to positive edges are I and whose blocks corresponding to negative edges are $-I$. The nonlinear Laplacian for this edge potential is the matrix $L_{\mathcal{F}}^S = \delta^T S \delta$, so it is in fact a linear operator.¹ These dynamics are more akin to those considered in [4].

The kernel of $L_{\mathcal{F}}^S$ contains $H^0(G; \mathcal{F})$, but may be larger. Further, this signed sheaf Laplacian is not necessarily positive semidefinite, so its corresponding diffusion dynamics may be unstable. One simple situation in which this happens is as follows:

PROPOSITION 12.1. *Suppose E_- is a cutset of G , dividing the graph into subgraphs G_0 and G_1 . If the natural map $H^0(G, G_1; \mathcal{F}) \rightarrow H^0(G_0; \mathcal{F})$ is not surjective — that is, if there exists a local section on G_0 which does not extend by zero to a global section of G — then $L_{\mathcal{F}}^S$ is indefinite.*

Proof. Take some nonzero $x \in H^0(G_0; \mathcal{F})$ which is not in the image of this map. Extending x by zero to the rest of G , we have $\delta_e x \neq 0$ for some $e \in E_-$, but $\delta_e x = 0$ for all $e \in E_+$. Thus, $\langle x, L_{\mathcal{F}}^S x \rangle = \sum_{e \in E_-} \langle \delta_e x, -\delta_e x \rangle$. This is negative because there is at least one nonzero term. \square

For the case of the constant sheaf on G , the relevant map is never surjective, so a graph with a cutset of negative edges always has an indefinite signed Laplacian, and hence unstable opinion dynamics.

13. Conclusions. This work introduced a novel and compelling ensemble of techniques from cellular sheaves, sheaf cohomology, and sheaf Laplacians, to model and analyze opinion dynamics over networks. The increase in mathematical sophistication comes with an increase in expressiveness of the model: private opinions, personalized expressions of opinions, evolution of communication structures, stubbornness, obfuscation, and bounded confidence are all easily expressed using discourse sheaves and sheaf diffusion dynamics. Despite this, there is no increase in difficulty of computation or analysis. The language of harmonic extension converts subtle solvability conditions to simple linear-algebraic cohomology computations. Despite the imposing terminology, sheaf cohomology is a concise and computable tool for determining feasibility of solving problems of existence, extension, approximation, controllability, and observability on sheaf dynamics.

The diffusive dynamics studied here have been linear or near-linear. Deeper analysis of the nonlinear dynamical systems (9.1) and (11.4) is needed, as well as the many

¹As is common in mathematics, “nonlinear” here really means “not-necessarily-linear.”

other possible extensions. These include incorporation of probabilistic elements or including multi-agent interactions using sheaves on simplicial complexes.

Since they are defined in abstract structural terms, sheaf dynamics are applicable to more than simply opinion dynamics. This paper may be regarded as an introduction to the dynamics of cellular sheaves as a broad generalization of network dynamics, with social networks and opinion dynamics serving as a running example. Other examples of networked dynamics effected by local evolution operators on rich data structures may be found in, e.g., neuroscience and game theory, at least. For the sake of accessibility, much of the language and techniques of sheaf theory and sheaf cohomology have been excised from this paper. A full incorporation of sheaf-theoretic operations would increase the precision, concision, and generality of this work.

REFERENCES

- [1] R. P. ABELSON, *Mathematical models of the distribution of attitudes under controversy*, in Contributions to Mathematical Psychology, N. Fredericksen and H. Gullicksen, eds., Holt, Rhinehart and Winston, 1964.
- [2] C. ALTAFINI, *Dynamics of opinion forming in structurally balanced social networks*, PLOS ONE, 7 (2012), p. e38135.
- [3] C. ALTAFINI, *Consensus problems on networks with antagonistic interactions*, IEEE Transactions on Automatic Control, 58 (2013), pp. 935–946.
- [4] C. ALTAFINI AND G. LINI, *Predictable dynamics of opinion forming for networks with antagonistic interactions*, IEEE Transactions on Automatic Control, 60 (2015), pp. 342–357.
- [5] M. BELKIN AND P. NIYOGI, *Laplacian eigenmaps for dimensionality reduction and data representation*, Neural Computation, 15 (2003), pp. 1373–1396.
- [6] V. D. BLONDEL, J. M. HENDRICKX, AND J. N. TSITSIKLIS, *Continuous-time average-preserving opinion dynamics with opinion-dependent communications*, SIAM Journal on Control and Optimization, 48 (2010), pp. 5214–5240.
- [7] C. CANUTO, F. FAGNANI, AND P. TILLI, *An Eulerian approach to the analysis of Krause’s consensus models*, SIAM Journal on Control and Optimization, 50 (2012), pp. 243–265.
- [8] C. CASTELLANO, S. FORTUNATO, AND V. LORETO, *Statistical physics of social dynamics*, Reviews of Modern Physics, 81 (2009), pp. 591–646.
- [9] F. CERAGIOLI AND P. FRASCA, *Continuous-time discontinuous equations in bounded confidence opinion dynamics*, IFAC Proceedings Volumes, 44 (2011), pp. 1986–1990.
- [10] F. CERAGIOLI AND P. FRASCA, *Continuous and discontinuous opinion dynamics with bounded confidence*, Nonlinear Analysis: Real World Applications, 13 (2012), pp. 1239–1251.
- [11] F. CHUNG, *Spectral Graph Theory*, AMS, 1992.
- [12] R. R. COIFMAN AND S. LAFON, *Diffusion maps*, Applied and Computational Harmonic Analysis, 21 (2006), pp. 5–30.
- [13] J. CURRY, *Sheaves, Cosheaves, and Applications*, PhD thesis, University of Pennsylvania, 2014.
- [14] G. DEFFUANT, D. NEAU, F. AMBLARD, AND G. WEISBUCH, *Mixing beliefs among interacting agents*, Advances in Complex Systems, 03 (2000), pp. 87–98.
- [15] M. H. DEGROOT, *Reaching a consensus*, Journal of the American Statistical Association, 69 (1974), pp. 118–121.
- [16] A. DIMCA, *Sheaves in Topology*, Springer-Verlag, Berlin, 2004.
- [17] J. C. DITTMER, *Consensus formation under bounded confidence*, Nonlinear Analysis: Theory, Methods & Applications, 47 (2001), pp. 4615–4621.
- [18] J. R. P. FRENCH, *A formal theory of social power*, Psychological Review, 63 (1956), pp. 181–194.
- [19] N. E. FRIEDKIN, *The Problem of Social Control and Coordination of Complex Systems in Sociology: A Look at the Community Cleavage Problem*, IEEE Control Systems Magazine, 35 (2015), pp. 40–51.
- [20] N. E. FRIEDKIN AND E. C. JOHNSEN, *Social influence and opinions*, The Journal of Mathematical Sociology, 15 (1990), pp. 193–206.
- [21] N. E. FRIEDKIN AND E. C. JOHNSEN, *Social influence networks and opinion change*, Advances in Group Processes, 16 (1999), pp. 1–29.
- [22] S. I. GELFAND AND Y. I. MANIN, *Methods of Homological Algebra*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, second edition ed., 2003.
- [23] J. HANSEN AND R. GHRIST, *Toward a spectral theory of cellular sheaves*, Journal of Applied

- and Computational Topology, 3 (2019), pp. 315–358.
- [24] R. HARTSHORNE, *Algebraic Geometry*, vol. 52 of Graduate Texts in Mathematics, Springer New York, New York, NY, 1977.
 - [25] R. HEGSELMANN AND U. KRAUSE, *Opinion dynamics and bounded confidence models, analysis and simulation*, Journal of Artificial Societies and Social Simulation, 5 (2002).
 - [26] A. JADBABAIE, J. LIN, AND A. S. MORSE, *Coordination of groups of mobile autonomous agents using nearest neighbor rules*, IEEE Transactions on Automatic Control, 48 (2003), p. 16.
 - [27] M. KASHIWARA AND P. SCHAPIRA, *Sheaves on Manifolds*, no. 292 in Grundlehren Der Mathematischen Wissenschaften, Springer-Verlag Berlin Heidelberg, 1990.
 - [28] A. KELLEY, *Stability of the center-stable manifold*, Journal of Mathematical Analysis and Applications, 18 (1967), pp. 336–344.
 - [29] K. LEHRER, *When Rational Disagreement is Impossible*, Noûs, 10 (1976), pp. 327–332.
 - [30] L. I. NICOLAESCU, *Lectures on the Geometry of Manifolds*, World Scientific, second edition ed., Sept. 2007.
 - [31] A. V. PROSKURNIKOV AND R. TEMPO, *A tutorial on modeling and analysis of dynamic social networks. Part I*, Annual Reviews in Control, 43 (2017), pp. 65–79.
 - [32] A. V. PROSKURNIKOV AND R. TEMPO, *A tutorial on modeling and analysis of dynamic social networks. Part II*, Annual Reviews in Control, 45 (2018), pp. 166–190.
 - [33] E. D. SONTAG, *Mathematical Control Theory: Deterministic Finite Dimensional Systems*, Texts in Applied Mathematics, Springer-Verlag, New York, second ed., 1998.
 - [34] K. SZNAJD-WERON AND J. SZNAJD, *Opinion evolution in closed community*, International Journal of Modern Physics C, 11 (2000), pp. 1157–1165.
 - [35] H. G. TANNER, A. JADBABAIE, AND G. J. PAPPAS, *Stable flocking of mobile agents, part I: Fixed topology*, in 42nd IEEE International Conference on Decision and Control, vol. 2, Dec. 2003, pp. 2010–2015 Vol.2.
 - [36] M. TAYLOR, *Towards a mathematical theory of influence and attitude change*, Human Relations, 21 (1968), pp. 121–139.
 - [37] W. WEIDLICH, *The statistical description of polarization phenomena in society*, British Journal of Mathematical and Statistical Psychology, 24 (1971), pp. 251–266.