Signed Opinion Dynamics with Information: Structure, Antagonism, and Long-Run Behavior

MMSS Thesis

Sophia Pi
Advised by Prof. Ben Golub

May 8, 2026

Contents

1 Introduction **3**

2 Models of Signed Opinion Dynamics **5**
2.1 Antagonistic Networks 5
2.2 Update Rules 5
2.3 External Information 7
2.4 Structural Balance 7

3 Existing Literature and Open Regimes **9**

4 Opposing Dynamics **11**
4.1 Without Information: Decay to Zero when Balance Fails 11
4.2 With Structural Balance: Misunderstood DeGroot 12
4.3 With Information: Convergence Without Structural Balance 14

5 Repelling Dynamics with Constant Information **16**
5.1 Model with Media Leakage 16
5.2 Spectral Monotonicity and the Lower Edge 17
5.3 Long-Run Fixed Point and The Stability Threshold 18

6 Time-Varying Information and Frequency Response **21**
6.1 Periodic Signals 21
6.2 Periodic Steady States 22
6.3 Gain Along Eigenmodes and Arbitrary Opinion Profiles 23
6.4 Effect of Network Properties and Antagonism on Gain 26

7 Optimizing Phase-Lagged Signals **28**
7.1 A Solved Analytical Example: The Frustrated Triangle 28
7.2 Numerical Methodology 31
7.3 Numerical Experiments 33
7.3.1 Three-Block Signed SBM 33

7.3.2 All-Antagonistic k-Cycles 34
7.3.3 Signed Lattices 34

8 Discussion and Open Questions **35**

A Spectral Tools **39**
A.1 Laplacian Preliminaries 39
A.2 Structural Balance and the Operator Norm 39
A.3 Boundedness and the Spectral Radius 40
A.4 Row-Sum and Spectral-Radius Regimes 40

B Doubly Infinite Sequence-Space Model **42**
B.1 Reduction to the forward constant-signal model 42
B.2 Existence of bounded solution histories 43

1 Introduction

Online disagreement often feels less like a failure to share information than a failure to share trust. On platforms such as Twitter, people do not merely encounter claims and update toward or away from them in isolation; they also react to who said them, which group appears to endorse them, and whether the source is perceived as friendly or hostile. Empirical work on online diffusion has found that false news spreads differently from true news on Twitter (Vosoughi et al., 2018), and that moral-emotional language shapes the spread of political content in social networks (Brady et al., 2017). These observations suggest that opinion dynamics models should account not only for information transmission, but also for antagonistic social relations.

Classical models of opinion dynamics usually begin from a more cooperative premise. In the DeGroot model (DeGroot, 1974), agents repeatedly average their neighbors' opinions using a nonnegative row-stochastic matrix. Under standard connectivity assumptions, the population converges to consensus. This framework is analytically clean and remains a natural benchmark. However, it has difficulty representing settings in which some social ties are oppositional: political opponents, rival firms, hostile media sources, or online communities that reject claims partly because they come from an opposing camp.

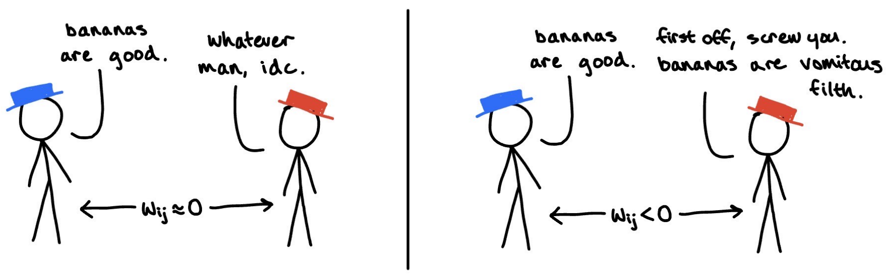

Figure 1: In nonnegative networks, the most an agent can do is disregard another agent's opinion. In antagonistic networks, an agent's opinion can be actively driven away from other opinions.

Signed-network models address this by allowing edges to carry positive or negative signs. Positive edges represent attraction, trust, or agreement; negative edges represent antagonism, distrust, or opposition. Two update rules are especially important. Under an **opposing** rule, an agent treats an enemy's opinion as something to be inverted before averaging. Under a **repelling** rule, an enemy's opinion pushes the agent away from that opinion. Although both rules use negative edges, they encode different behavioral mechanisms.

This thesis studies what happens when signed opinion dynamics are combined with external information. Two papers are the primary antecedents. Shi et al. (2019) study opposing and repelling dynamics on signed networks, but they do not include external information. Lena et al. (2023) introduce external information, but they study the opposing model with constant information and assume structural balance throughout. These papers leave open the regime that is the main focus of this thesis: repelling signed dynamics with external information, especially when that information varies over time.

The first part of the thesis revisits opposing dynamics. This is not the main new model, but it clarifies what the existing opposing framework can and cannot express. Without information, unbalanced opposing dynamics collapse to zero under standard assumptions. With structural balance, opposing dynamics reduce exactly to ordinary DeGroot averaging after flipping the sign convention of one camp. In this sense, the balanced opposing model is closer to "misunderstood DeGroot" than to a model of genuine repulsion. With external information, convergence no longer requires structural balance: positive information weights themselves make the social update contractive.

The second part introduces repelling dynamics with constant external information. Starting from the Shi repelling model, we add a media-leakage term: each agent places weight c_i on an external information source, with that weight subtracted from her self-retention term. The resulting update matrix is

$$M'(\beta) = I - \alpha L_{G^+} + \beta L_- - D,$$

where α controls attraction along positive ties, β controls repulsion along negative ties, and $D = \mathrm{diag}(c)$ records information weights. We then derive a stability threshold in β such that below this threshold, opinions converge to a unique fixed point, while above it, the system generically diverges.

The third part moves from constant information to oscillating information. Constant-information models can ask whether opinions converge to a fixed point or become unstable. They cannot describe persistent cycles in external signals, such as recurring news attention, market sentiment, or social-media outrage cycles. We therefore study signals of the form

$$\theta(t) = \Re(ve^{i\omega t}).$$

In the stable regime, we derive the periodic steady-state response that such inputs generate, governed by a frequency-response matrix. We also prove modal gain formula that characterizes the amplification of the opinion vector (in the long run) with respect to certain vectors (which may be interpreted as opinion profiles).

The stable repelling model therefore admits a frequency-response interpretation. The underlying linear-system calculation is general, but the signed-network content enters through how the antagonism parameter shifts the spectrum and changes which modes are amplified.

The final part studies phase-lagged information. If agents receive the same oscillating signal at different magnitudes and lags, then the signal can be written as $a_i = r_i e^{i\phi_i}$. We ask which phase-lagged signal maximizes disagreement, measured by time-averaged variance of the real steady-state trajectory. We analytically solve the toy example of the fully negative ("frustrated") triangle, where we show that the optimal lags are not binary $0/\pi$. Instead, the three agents are assigned phases separated by 120°. Numerical experiments suggest that similar multi-spoke phase patterns appear in larger frustrated signed networks, including three-block signed stochastic block models, all-antagonistic k-cycles, and signed lattices.

The thesis is organized as follows. Section 2 introduces signed networks, update rules, external information, and structural balance. Section 3 situates the thesis relative to existing work and identifies the open regimes. Section 4 revisits opposing dynamics and clarifies the role of structural balance. Section 5 develops the repelling model with constant information and proves the stability threshold. Section 6 introduces oscillating information and derives the frequency-response theory. Section 7 studies optimized phase-lagged signals and reports numerical experiments. Section 8 discusses open questions. Additional spectral tools and the doubly infinite sequence-space formulation are collected in the appendices.

2 Models of Signed Opinion Dynamics

In this section, we introduce the primary models of signed opinion dynamics/antagonistic networks and describe our working concepts of structural balance and (constant) external information.

2.1 Antagonistic Networks

Let $G = (V, E)$ be a signed graph on n agents, $V = \{1, \ldots, n\}$, with a partition of edges into positive (friendly) edges E^+ and negative (antagonistic) edges E^-. For each agent i, write $N_i^+ = \{j : \{i,j\} \in E^+\}$ and $N_i^- = \{j : \{i,j\} \in E^-\}$ for the sets of friends and opponents, with degrees $\deg_i^+ = |N_i^+|$ and $\deg_i^- = |N_i^-|$.

The general language of signed networks allows directed and heterogeneous weights. However, the main spectral results in Sections 5 and 6 assume undirected signed graphs with symmetric update matrices. Unless otherwise stated, G^+ and G^- are undirected, L_{G^+} and L_- are symmetric graph Laplacians, and $D = \mathrm{diag}(c)$ with $c_i > 0$.

Define the following matrices. Let A_{G^+} be the (unsigned) adjacency matrix of G^+ (with $[A_{G^+}]_{ij} = 1$ if $\{i,j\} \in E^+$), and A_{G^-} similarly for G^-. Let $D_{G^+} = \mathrm{diag}(\deg_1^+, \ldots, \deg_n^+)$ and $D_{G^-} = \mathrm{diag}(\deg_1^-, \ldots, \deg_n^-)$ be the corresponding degree matrices. The standard Laplacians are:

$$L_{G^+} := D_{G^+} - A_{G^+}, \qquad L_- := D_{G^-} - A_{G^-}.$$

Both L_{G^+} and L_- are positive semidefinite (PSD); see Theorem A.1.

For the repelling model, we also use the **antagonistic Laplacian**

$$L^r_{G^-} := -D_{G^-} + A_{G^-} = -L_-,$$

which is negative semidefinite (NSD). As a standalone matrix, $L^r_{G^-}$ has nonpositive diagonal entries and nonnegative off-diagonal entries. In the repelling update, however, it appears with a negative sign, so $-\beta L^r_{G^-} = \beta L_-$. Thus the final update matrix receives a positive diagonal contribution from antagonistic degree and negative off-diagonal contributions from enemies.

2.2 Update Rules

We study the following family of linear update rules. Opinions at time t are collected in a vector $x(t) \in \mathbb{R}^n$. The general form is:

$$x(t+1) = \tilde{W}\, x(t) + c \odot \theta(t), \tag{1}$$

where $\tilde{W}$ encodes the social interaction (including signs), $c \in \mathbb{R}^n$ is a vector of media-attention weights with $c_i \geq 0$, and $\theta(t) \in \mathbb{R}^n$ is the vector of external information ("media sentiments") at time t.

The two canonical rules implemented by $\tilde{W}$ are:

- **Opposing rule:**
$$x_i(t+1) = \sum_{j \in N_i^+} w_{ij}\, x_j(t) \;-\; \sum_{j \in N_i^-} w_{ij}\, x_j(t).$$
  Agent i averages friends' opinions and subtracts (flips and averages) enemies' opinions. In matrix form, $\tilde{W} = W \odot S$ where $W = |\tilde{W}|$ is entrywise nonneg and row-stochastic, and $S_{ij} = +1$ if i and j are friends, $S_{ij} = -1$ if they are opponents.

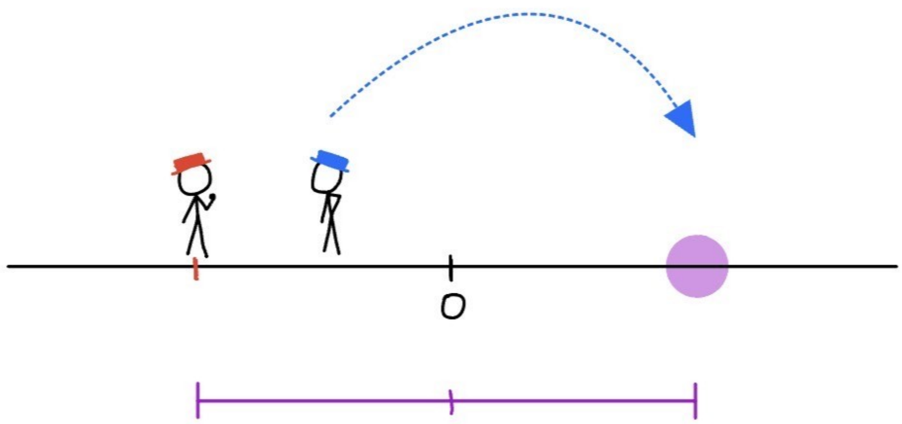

Figure 2: Under opposing dynamics, agents are pulled towards the inverse (about 0) of their enemies' opinions.

- **Repelling rule:**

$$x_i(t+1) = \left(1 - \alpha \deg_i^+ + \beta \deg_i^-\right) x_i(t) + \alpha \sum_{j \in N_i^+} x_j(t) - \beta \sum_{j \in N_i^-} x_j(t).$$

Equivalently,

$$x_i(t+1) = x_i(t) - \alpha \sum_{j \in N_i^+} (x_i(t) - x_j(t)) + \beta \sum_{j \in N_i^-} (x_i(t) - x_j(t)).$$

Here $\alpha \in [0,1]$ is a friendliness weight and $\beta \in [0,1]$ an antagonism weight. In matrix form,

$$M(\beta) := I - \alpha L_{G^+} - \beta L^r_{G^-} = I - \alpha L_{G^+} + \beta L_-. \tag{2}$$

Note that $M(\beta)\mathbf{1} = \mathbf{1}$ (row sums equal 1). More details are discussed in section 5.1.

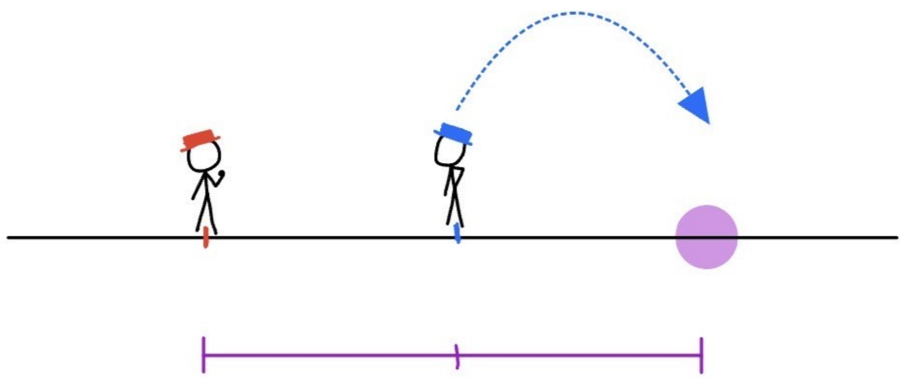

Figure 3: Under repelling dynamics, agents are pushed away from their enemies' opinions, independent of absolute points like 0.

The opposing rule models agents who "disagree with their enemies," while the repelling rule models agents who actively move away from enemies' positions.

2.3 External Information

The term $c \odot \theta(t)$ in (1) captures external information entry. We consider two regimes:

- **Constant (time-invariant) information:** $\theta(t) = \theta$ for all t. This is the case analyzed by Lena et al. (2023).
- **Time-varying information:** $\theta(t)$ is a sequence (e.g., drifting, periodic, or group-specific). This case is treated in Section 6.

When information is constant, the update rule (1) defines an affine dynamical system $x(t+1) = \tilde{W}\, x(t) + u$, where $u = c \odot \theta$ is a fixed input vector.

2.4 Structural Balance

Structural balance corresponds to the colloquial principle "enemy of my enemy is my friend." When structural balance holds, there exists a **signature matrix** $D_S = \mathrm{diag}(\pm 1)$ such that $D_S \tilde{W} D_S^{-1} \geq 0$ entrywise, a diagonal similarity that transforms the signed problem into an unsigned one.

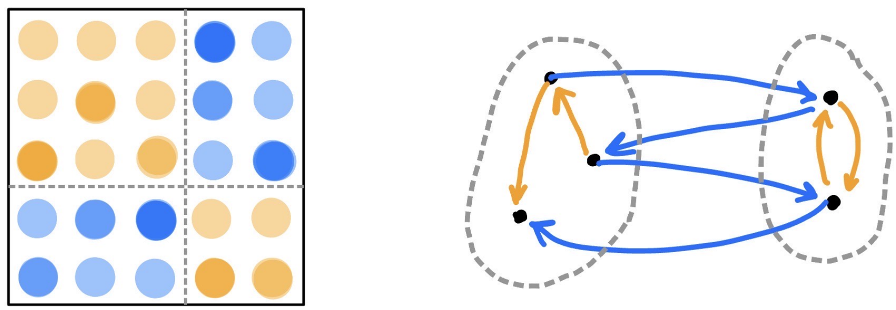

Figure 4: In structurally balanced networks, agents can be cleanly separated into two camps that are "friendly" internally and antagonistic between camps.

Definition 2.1 (Structural Balance)**.** A signed graph G is **structurally balanced** if every cycle in G has nonnegative sign product (i.e., an even number of negative edges). Equivalently (Cartwright and Harary, 1956), the vertex set V can be partitioned into two groups A and B such that all edges within each group are positive and all edges between groups are negative.

When structural balance fails, no such similarity exists, and the dynamics can behave qualitatively differently.

3 Existing Literature and Open Regimes

This thesis is primarily motivated by two papers. The first is Shi et al. (2019), which studies signed-network dynamics without external information. The second is Lena et al. (2023), which studies opposing dynamics with constant external information. Together, these papers give a natural map of the existing literature: they distinguish opposing from repelling rules, and they distinguish models with and without external information. They also leave open the regime that is central to this thesis: repelling signed dynamics with information.

Shi et al. (2019) provide a comprehensive study of opinion dynamics over signed networks without external information. They distinguish two update rules. Under the **opposing** rule, a negative neighbor's opinion is inverted before entering the update. Under the **repelling** rule, negative neighbors push agents away from their current opinions. This distinction is important: opposing dynamics are closely tied to structural balance and bipartite consensus, while repelling dynamics can produce stability thresholds and divergence. Their analysis, however, is carried out under a restrictive $I -$ Laplacian form of the update matrix, which enforces symmetry and limits weight heterogeneity.

Lena et al. (2023) study a different extension: opposing dynamics with constant external information. In their model, agents place weight on both social neighbors and external media sentiments. They derive closed-form long-run opinions in terms of signed Katz–Bonacich centralities. Their analysis assumes structural balance throughout. One observation of this thesis, developed in Section 4.3, is that structural balance is not needed for the basic convergence result once every agent places positive weight on external information. It remains important for interpreting the fixed point, but not for existence or convergence itself.

The existing work naturally suggests several questions. First, what exactly is structural balance doing in opposing dynamics? Second, if external information stabilizes opposing dynamics, can we add information to the repelling model as well? Third, once information is included, should we restrict attention to constant information, or can signed networks be used to study genuinely time-varying signals?

Tables 1 and 2 summarize the regimes that are already understood and the regimes this thesis addresses. We discuss the relevant parts of Shi et al. (2019) and Lena et al. (2023) in more detail in the sections where they become mathematically useful: opposing dynamics in Section 4, repelling dynamics in Section 5, and oscillating information in Section 6.

Tables 1 and 2 summarize the regimes that are already understood and the regimes this thesis addresses.

Table 1: **No external information.** Summary of known results and open questions.

	Opposing	**Repelling**
SB	✓$_*$ ✓ Bipartite consensus	✓$_*$ Consensus ($\beta < \beta^*$); divergence ($\beta > \beta^*$)
No SB	✓$_*$ Decay to 0	✓$_*$ Same thresholds (SB irrelevant for repelling)

Legend: ■ = Shi et al. (2019); ■ = Lena et al. (2023); $_*$ = restrictive $I -$ Laplacian assumption.

The first table shows that, without information, opposing and repelling dynamics are already

Table 2: **With external information (constant θ).** Contributions of this thesis.

	Opposing	**Repelling**
SB	✓ Bipartite consensus (constant info)	? this thesis
No SB	? this thesis	? this thesis

Legend: ? = addressed in this thesis.

well studied. However, it also suggests why opposing dynamics may be limited as a model of antagonistic amplification. Without structural balance, the opposing model decays to zero; with structural balance, it behaves like DeGroot after a sign change. This motivates Section 4, where we clarify the role of structural balance in the opposing model.

The second table shows the main opening. Once external information is introduced, the opposing model has been studied, but the analogous repelling model has not. This is the starting point of 5: we add media weights to the Shi repelling model and derive a stability threshold for the antagonism parameter β. The threshold compares the destabilizing force of negative ties against the stabilizing force of positive ties and external information.

Finally, the tables only describe constant-information regimes. A central point of this thesis is that constant information is not the only natural informational environment. Many real signals vary over time: news cycles, market sentiment, and social-media attention all rise and fall. After establishing the constant-information repelling model, we therefore turn to oscillating information. This leads to the frequency-response analysis in Section 6 and the phase-lag optimization problem studied later.

4 Opposing Dynamics

Before turning to repelling dynamics, we briefly revisit the opposing model studied in Lena et al. (2023). This model is the closest existing signed-network model with external information, so it provides a useful point of comparison. The observations in this section are meant to clarify which parts of the model rely on structural balance and which parts rely only on the contraction created by positive information weights.

Conceptually, opposing dynamics treat a negative neighbor's opinion as something to be inverted. If i dislikes j, then j's positive opinion enters i's update negatively. This is different from the repelling model studied later, where negative ties cause agents to move away from one another. This distinction matters: we observe that opposing dynamics (particularly in the informationless case) are closely related to DeGroot averaging after an appropriate sign change (which we call "misunderstood DeGroot"), while repelling dynamics can generate amplification and instability behavior that this opposing model cannot.

4.1 Without Information: Decay to Zero when Balance Fails

We first consider the no-information opposing model $x(t+1) = \tilde{W}x(t)$. As proven in Lena et al. (2023), the structurally balanced case produces bipartite consensus. A natural complementary question, then, is what happens when structural balance fails. Under the usual self-confidence and irreducibility assumptions, we show that the answer turns out to be somewhat anticlimactic: there is no nonzero long-run opinion pattern.

The lynchpin in the argument is the following lemma, which gives the straightforward spectral reason for this "vanishing to zero" behavior. It slightly generalizes the analogous observation in Shi et al. (2019), while keeping the same interpretation: without structural balance, the opposing model loses its peripheral eigenvalues and becomes a contraction. The proof uses the same basic principle highlighted in Theorem A.3: when structural balance holds, the signs of $\tilde{W}$ can be removed by a diagonal sign change.

Lemma 4.1 (No unit eigenvalue without structural balance)**.** *Let $\tilde{W} \in \mathbb{R}^{n\times n}$, and let $W = |\tilde{W}|$. Suppose W is row-stochastic and irreducible, and suppose agents have positive self-weight, $w_{ii} > 0$ for all i. If $\tilde{W}$ is not structurally balanced, then no eigenvalue of $\tilde{W}$ has magnitude exactly equal to 1.*

Proof. We give the argument because it explains the role of structural balance. Suppose $\tilde{W}y = \lambda y$ for some nonzero y with $|\lambda| = 1$. Taking absolute values row by row gives

$$|y_i| = \Big|\sum_j \tilde{w}_{ij}y_j\Big| \le \sum_j |\tilde{w}_{ij}|\,|y_j| = (W|y|)_i.$$

Thus $W|y| \ge |y|$ entrywise. Let $\varepsilon \gg 0$ be the left Perron eigenvector of W. If the inequality were strict in any coordinate, then

$$\varepsilon^\top W|y| > \varepsilon^\top |y|,$$

contradicting $\varepsilon^\top W = \varepsilon^\top$. Hence $W|y| = |y|$. Since W is irreducible and row-stochastic, $|y|$ is proportional to $\mathbf{1}$. After rescaling, assume $|y_i| = 1$ for all i.

Equality must therefore hold in the triangle inequality for every row:

$$1 = \Big|\sum_j \tilde{w}_{ij} y_j\Big| = \sum_j |\tilde{w}_{ij}|\,|y_j|.$$

Hence, for each i, all nonzero terms $\tilde{w}_{ij} y_j$ have the same complex argument. Writing $d_i := y_i/|y_i|$, this implies that along every nonzero edge,

$$\tilde{w}_{ij} = \lambda d_i w_{ij} d_j^{-1}.$$

The positive self-weight assumption gives $\tilde{w}_{ii} = w_{ii} > 0$, so the same identity with $i = j$ forces $\lambda = 1$. Since $\tilde{W}$ is real and $w_{ij} \geq 0$, the factors $d_i d_j^{-1}$ must be real signs on every edge. Thus there exists a diagonal sign matrix S such that

$$S\tilde{W}S^{-1} = W \geq 0.$$

This is exactly the diagonal sign-change condition defining structural balance, and is the same mechanism used in Lemma A.3. This contradicts the assumption. Therefore $\tilde{W}$ has no unit-modulus eigenvalue. □

Corollary 4.2 (Decay to zero under opposing dynamics without balance)**.** *Consider opposing dynamics without external information,*

$$x(t+1) = \tilde{W}x(t),$$

with $W = |\tilde{W}|$ row-stochastic and irreducible, and with $w_{ii} > 0$ for all i. If $\tilde{W}$ is not structurally balanced, then $x(t) \to 0$ as $t \to \infty$ for every initial condition $x(0)$.

Proof. Since $W = |\tilde{W}|$ is row-stochastic, $\rho(\tilde{W}) \leq \rho(W) = 1$. This is the boundary case of the absolute row-sum control summarized in Fact A.5 and Corollary A.6. By Lemma 4.1, no eigenvalue of $\tilde{W}$ has magnitude 1. Hence $\rho(\tilde{W}) < 1$, so

$$x(t) = \tilde{W}^t x(0) \to 0.$$

□

Remark 4.3. This corollary identifies a limitation of the no-information opposing rule. Structural balance is not merely a convenient sufficient condition for bipartite consensus; under the assumptions above, it is the condition that prevents all opinions from decaying to zero. Thus, without external information, the opposing model has two sharply different regimes: balanced networks behave like a signed version of DeGroot learning, while unbalanced networks contract.

4.2 With Structural Balance: Misunderstood DeGroot

We now make an important conceptual claim. The structurally balanced case is not merely a technically convenient regime in which opposing dynamics behave well. It also reveals something conceptually important about the model. When structural balance holds, the apparent antagonism in the opposing rule can be removed by a change of signs: one camp says “positive” where the other says “negative,” but after “translating” between their labels, both camps are in fact performing ordinary DeGroot averaging (or at the very least, the updating mechanisms are mathematically identical).

In this sense, we argue that the opposing model under structural balance is closer to a model of miscommunication, or of two groups speaking different sign conventions, than to a model of genuine antagonism. We state this precisely in the following proposition.

Proposition 4.4 (LMZ reduces to DeGroot)**.** *Under opposing dynamics with structural balance, let $V = A \cup B$, and define the group-identity vector $\mathbf{1}_-$ by*

$$(\mathbf{1}_-)_i = \begin{cases} +1, & i \in A, \\ -1, & i \in B. \end{cases}$$

Then the following "misunderstood DeGroot" algorithm produces the same trajectory as the opposing model:

1. *Translate: $\hat{x}(0) = \mathbf{1}_- \odot x(0)$.*
2. *Run DeGroot: $\hat{x}(t) = W^t \hat{x}(0)$, where $W = |\tilde{W}|$.*
3. *Translate back: $x_i(t) = (\mathbf{1}_-)_i \hat{x}_i(t)$.*

Equivalently, for all $t \geq 0$,

$$x(t) = \mathbf{1}_- \odot \hat{x}(t).$$

Proof. The proof is by induction. The base case follows from the definition of $\hat{x}(0)$. Suppose $\hat{x}(t) = \mathbf{1}_- \odot x(t)$.

For $i \in A$, positive edges go to A and negative edges go to B. Therefore

$$\hat{x}_i(t+1) = (W\hat{x}(t))_i = \sum_{j \in A} w_{ij} x_j(t) + \sum_{j \in B} w_{ij}(-x_j(t)) = (\tilde{W}x(t))_i = x_i(t+1).$$

For $i \in B$, we have $\hat{x}_j(t) = x_j(t)$ for $j \in A$ and $\hat{x}_j(t) = -x_j(t)$ for $j \in B$. Therefore

$$\hat{x}_i(t+1) = \sum_{j \in A} w_{ij} x_j(t) - \sum_{j \in B} w_{ij} x_j(t).$$

Since edges from B to A are negative and edges within B are positive,

$$(\tilde{W}x(t))_i = -\sum_{j \in A} w_{ij} x_j(t) + \sum_{j \in B} w_{ij} x_j(t).$$

Thus

$$\hat{x}_i(t+1) = -(\tilde{W}x(t))_i = -x_i(t+1).$$

So $\hat{x}_i(t+1) = (\mathbf{1}_-)_i x_i(t+1)$ for every i, completing the induction. □

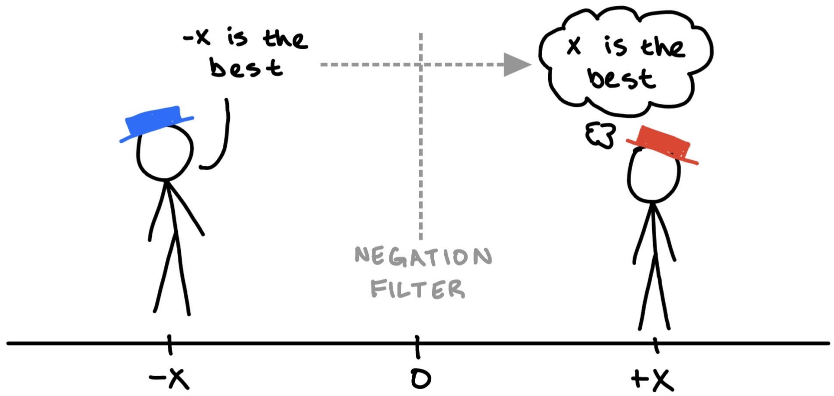

Figure 5: Enter Caption

This reduction gives a useful interpretation of bipartite consensus. The two camps do not converge to unrelated values. Rather, after flipping the signs of one camp's opinions, the system is exactly standard DeGroot learning. If $\hat{x}(t)$ converges to a scalar consensus $\hat{c}$, then agents in A converge to $\hat{c}$, while agents in B converge to $-\hat{c}$.

4.3 With Information: Convergence Without Structural Balance

The informational model of Lena et al. (2023) is more interesting because external information prevents the no-information collapse described above. In their setting, agents update according to

$$x(t+1) = \tilde{W}x(t) + c \odot \theta, \tag{3}$$

where $\theta \in \mathbb{R}^n$ is a constant media sentiment vector and $c_i > 0$ is the weight agent i places on media. In LMZ, structural balance is maintained throughout the analysis, which allows the fixed point to be interpreted through signed Katz–Bonacich centralities.

For convergence itself, however, structural balance is stronger than necessary. The contraction comes from the fact that positive media weights make the social part of the update absolutely sub-row-stochastic, so the spectral-radius test in Fact A.5 and Corollary A.6 applies directly.

Proposition 4.5 (Convergence in the informational opposing model does not require balance)**.** *Suppose $\tilde{W} \in \mathbb{R}^{n\times n}$, let $W = |\tilde{W}|$, and assume*

$$\sum_j w_{ij} = 1 - c_i \qquad \text{with} \qquad c_i > 0$$

for every i. Then

$$\rho(\tilde{W}) < 1.$$

Consequently, for every constant $\theta \in \mathbb{R}^n$, the affine system

$$x(t+1) = \tilde{W}x(t) + c \odot \theta$$

has a unique globally attracting fixed point

$$x^* = (I - \tilde{W})^{-1}(c \odot \theta),$$

regardless of whether $\tilde{W}$ is structurally balanced.

Proof. Since $W = |\tilde{W}|$,

$$\|\tilde{W}\|_\infty = \max_i \sum_j |\tilde{w}_{ij}| = \max_i \sum_j w_{ij} = \max_i (1 - c_i).$$

Because $c_i > 0$ for every i, if $c_{\min} := \min_i c_i$, then

$$\|\tilde{W}\|_\infty = 1 - c_{\min} < 1.$$

Thus $\tilde{W}$ is absolutely sub-row-stochastic with strict row-sum bound. By Fact A.5 and Corollary A.6,

$$\rho(\tilde{W}) \leq \|\tilde{W}\|_\infty < 1.$$

It follows that $I - \tilde{W}$ is invertible and that

$$\sum_{t=0}^{\infty} \tilde{W}^t$$

converges. The fixed point is therefore

$$x^* = (I - \tilde{W})^{-1}(c \odot \theta).$$

For any trajectory, setting $y(t) = x(t) - x^*$ gives

$$y(t+1) = \tilde{W} y(t),$$

so $y(t) \to 0$. □

This proposition should be read as a modest clarification of the LMZ informational model. Structural balance remains important for interpreting the fixed point as bipartite agreement or disagreement, and for obtaining the particular signed centrality formulas emphasized in their paper. But it is not needed for the basic convergence claim once every agent places positive weight on external information. In the informational opposing model, the media weights themselves create the contraction.

The observation we make in this section are therefore threefold. First, without information, opposing dynamics need structural balance to avoid collapse to zero. Second, when structural balance is imposed, the model becomes much less antagonistic than it first appears: after flipping one camp's sign convention, the dynamics reduce exactly to ordinary DeGroot averaging. Thus the balanced opposing model is best understood as "misunderstood DeGroot," rather than as a model of genuine repulsion. Third, once external information is added, convergence no longer requires structural balance; positive media weights make the social update contractive even in unbalanced networks. These observations motivate the repelling model studied next, where negative ties cannot be simply removed by a change of language and can instead destabilize or amplify disagreement.

5 Repelling Dynamics with Constant Information

We now turn to the main model studied in this thesis: repelling signed dynamics with external information. The opposing model is useful as a first signed extension of DeGroot learning, but it has a limitation. When the network is structurally balanced, opposing dynamics can be interpreted as ordinary averaging after flipping the signs of one camp's opinions; when the network is not structurally balanced, the no-information model often collapses toward zero. Thus negative weights enter the model, but they do not yet produce a rich theory of persistent antagonistic amplification.

The repelling model of Shi et al. (2019) treats negative ties differently. Instead of being drawn towards a negative neighbor's inverted opinion, an agent moves away from (is repelled by) that neighbor's opinion. This makes antagonism a destabilizing force: strong negative ties can amplify disagreement rather than merely reverse signs. In this section, we characterize what happens when this repelling dynamic is combined with constant external information.

5.1 Model with Media Leakage

Recall from (2) that the repelling update matrix without information is

$$M(\beta) = I - \alpha L_{G^+} + \beta L_-, \qquad M(\beta)\mathbf{1} = \mathbf{1}.$$

Here α controls attraction along positive edges and β controls repulsion along negative edges. The positive-edge term $-\alpha L_{G^+}$ pulls friends together, while the negative-edge term $+\beta L_-$ pushes enemies apart.

We introduce external information through a media-leakage mechanism. Each agent i places weight $c_i > 0$ on an external information source with sentiment θ_i. In this specification, the media weight is subtracted from agent i's self-retention term rather than by rescaling the entire social row. Let $D = \operatorname{diag}(c)$. The resulting model is

$$x(t+1) = M'(\beta)x(t) + D\theta, \tag{4}$$

where

$$M'(\beta) := M(\beta) - D = I - \alpha L_{G^+} + \beta L_- - D. \tag{5}$$

Thus

$$M'(\beta)_{ii} = 1 - \alpha \deg_i^+ + \beta \deg_i^- - c_i,$$

and the ith row sum of $M'(\beta)$ is $1 - c_i < 1$. The lost self-weight is exactly the weight assigned to external information.

It is useful to isolate the terms that counteract repulsion. Define

$$A := \alpha L_{G^+} + D.$$

Then

$$I - M'(\beta) = A - \beta L_-.$$

The matrix A represents the stabilizing part of the model: positive ties penalize disagreement among friends, and D anchors agents to external information. The term βL_- represents the destabilizing part: negative ties push agents away from enemies.

Throughout this section, we assume that agents assign nonnegative weight to their own past opinions in the absence of antagonism:

$$1 - \alpha \deg_i^+ - c_i \geq 0 \qquad \text{for every } i.$$

Equivalently, $\alpha \deg_i^+ + c_i \leq 1$. This is a baseline admissibility condition: before repelling effects are added, the update rule should not require agents to put negative weight on their own previous opinions. A related boundedness condition appears in Shi et al. (2019), who assume $0 < \alpha + \beta < 1/\max_i \deg_i$ in their model.

5.2 Spectral Monotonicity and the Lower Edge

Our stability analysis relies on spectral observations. Since $M'(\beta)$ is symmetric, the system is stable exactly when all eigenvalues lie inside the unit circle. The next two lemmas show that, under the admissibility condition above, the lower edge of the spectrum does not bind, and increasing antagonism moves all eigenvalues weakly upward.

Lemma 5.1 (Lower spectral edge is harmless under nonnegative baseline self-weights)**.** *Let*

$$M'(\beta) = I - \alpha L_{G^+} + \beta L_- - D,$$

where G^+ and G^- are undirected graphs, L_{G^+} and L_- are their graph Laplacians, $D = \operatorname{diag}(c_1, \ldots, c_n)$, and $\alpha, \beta \geq 0$. Suppose that

$$c_i > 0 \qquad \textit{and} \qquad \alpha \deg_i^+ + c_i \leq 1$$

for every agent i. Then

$$\lambda_{\min}(M'(\beta)) > -1 \qquad \textit{for every } \beta \geq 0.$$

Proof. First consider

$$M'(0) = I - \alpha L_{G^+} - D.$$

Its diagonal entries are $1 - \alpha \deg_i^+ - c_i$, and its off-diagonal entries are α along positive edges and 0 otherwise. By assumption, $M'(0)$ is entrywise nonnegative. Its ith row sum is $1 - c_i$. Since $c_i > 0$ for all i, if $c_{\min} := \min_i c_i$, then

$$\|M'(0)\|_\infty = \max_i (1 - c_i) = 1 - c_{\min} < 1.$$

Therefore $\rho(M'(0)) \leq \|M'(0)\|_\infty < 1$. Since $M'(0)$ is symmetric, all eigenvalues are real, so $\lambda_{\min}(M'(0)) > -1$.

Now

$$M'(\beta) = M'(0) + \beta L_-.$$

Since $L_- \succeq 0$, the Courant–Fischer characterization gives

$$\lambda_{\min}(M'(\beta)) = \min_{\|x\|=1} \left(x^\top M'(0)x + \beta x^\top L_- x \right) \geq \min_{\|x\|=1} x^\top M'(0)x = \lambda_{\min}(M'(0)).$$

Hence $\lambda_{\min}(M'(\beta)) > -1$. $\square$

Lemma 5.2 (Eigenvalues are nondecreasing in β)**.** *For $\beta_1 < \beta_2$, $\lambda_k(M'(\beta_1)) \le \lambda_k(M'(\beta_2))$ for all $k = 1, \dots, n$, with eigenvalues ordered increasingly.*

Proof. We have
$$M'(\beta_2) - M'(\beta_1) = (\beta_2 - \beta_1)L_-.$$
Since $\beta_2 > \beta_1$ and $L_- \succeq 0$, this difference is positive semidefinite. Weyl monotonicity for Hermitian matrices therefore gives
$$\lambda_k(M'(\beta_1)) \le \lambda_k(M'(\beta_2))$$
for every k, with eigenvalues ordered increasingly. □

Together, these lemmas reduce the stability problem to a single question: when does the top eigenvalue of $M'(\beta)$ cross 1? In the following subsection, we derive the answer in the form of a threshold in the antagonism parameter β.

5.3 Long-Run Fixed Point and The Stability Threshold

We now ask when the system (4) converges. Since increasing β increases the eigenvalues of $M'(\beta)$, stronger antagonism eventually threatens stability. The threshold below identifies the exact point at which the destabilizing force of negative ties overwhelms the stabilizing force of positive ties and external information.

Theorem 5.3 (Stability threshold for repelling dynamics with information)**.** *Let G^+ and G^- be undirected graphs on n nodes. Let $\alpha \in [0, 1]$, $c \in \mathbb{R}^n$ with $c_i > 0$, $D = \mathrm{diag}(c)$, and $\theta \in \mathbb{R}^n$ constant. Define $M'(\beta)$ as in (5) and $A = \alpha L_{G^+} + D$. Since $c_i > 0$ for all i, $D \succ 0$, and hence $A \succ 0$. Define*
$$\beta_{\mathrm{thresh}}(c) := \inf_{\substack{x \neq 0 \\ x^\top L_- x > 0}} \frac{x^\top A x}{x^\top L_- x} = \inf_{\substack{x \neq 0 \\ x^\top L_- x > 0}} \frac{x^\top (\alpha L_{G^+} + D) x}{x^\top L_- x}. \tag{6}$$

Then:

1. ***Stable regime*** *($\beta < \beta_{\mathrm{thresh}}(c)$): $\lambda_{\max}(M'(\beta)) < 1$. In particular, $\rho(M'(\beta)) < 1$, and (4) converges to the unique fixed point*
$$x^* = (I - M'(\beta))^{-1} D\theta = (A - \beta L_-)^{-1} D\theta.$$
2. ***Unstable regime*** *($\beta > \beta_{\mathrm{thresh}}(c)$): $\lambda_{\max}(M'(\beta)) \ge 1$, and the system generically diverges.*

Proof. Since $M'(\beta)$ is real symmetric, it has an orthogonal diagonalization $M'(\beta) = Q\Lambda Q^\top$. Setting $z(t) = Q^\top x(t)$ and $b = Q^\top D\theta$, the update rule decouples into scalar recurrences
$$z_k(t+1) = \lambda_k z_k(t) + b_k, \qquad k = 1, \dots, n.$$
Each scalar recurrence converges if and only if $|\lambda_k| < 1$. Thus the full system converges if and only if $\rho(M'(\beta)) < 1$.

By Theorem 5.1, the lower spectral edge satisfies $\lambda_{\min}(M'(\beta)) > -1$. Hence stability reduces to

$$\lambda_{\max}(M'(\beta)) < 1.$$

Equivalently,

$$I - M'(\beta) = A - \beta L_- \succ 0.$$

First suppose $\beta < \beta_{\text{thresh}}(c)$. Take any nonzero x. If $x^\top L_- x = 0$, then $x \in \ker(L_-)$, so

$$x^\top (A - \beta L_-)x = x^\top A x > 0$$

because $A \succ 0$. If $x^\top L_- x > 0$, then by definition of $\beta_{\text{thresh}}(c)$,

$$\frac{x^\top A x}{x^\top L_- x} \geq \beta_{\text{thresh}}(c) > \beta.$$

Thus $x^\top A x > \beta x^\top L_- x$, so

$$x^\top (A - \beta L_-)x > 0.$$

Therefore $A - \beta L_- \succ 0$, and $\lambda_{\max}(M'(\beta)) < 1$.

In this stable regime, $I - M'(\beta)$ is invertible. Any fixed point satisfies

$$x^* = M'(\beta)x^* + D\theta,$$

so

$$x^* = (I - M'(\beta))^{-1} D\theta = (A - \beta L_-)^{-1} D\theta.$$

Moreover, if $y(t) = x(t) - x^*$, then $y(t+1) = M'(\beta)y(t)$, so $y(t) \to 0$ because $\rho(M'(\beta)) < 1$.

Now suppose $\beta > \beta_{\text{thresh}}(c)$. By definition of the infimum, there exists a nonzero $\hat{x}$ with $\hat{x}^\top L_- \hat{x} > 0$ such that

$$\hat{x}^\top A \hat{x} < \beta\, \hat{x}^\top L_- \hat{x}.$$

Thus

$$\hat{x}^\top (A - \beta L_-)\hat{x} < 0.$$

So $A - \beta L_-$ is not positive semidefinite, and therefore $\lambda_{\max}(M'(\beta)) \geq 1$. For initial conditions or inputs with nonzero projection onto an unstable eigendirection, the corresponding modal coordinate grows without bound. □

The threshold $\beta_{\text{thresh}}(c)$ has a direct behavioral interpretation. The denominator $x^\top L_- x$ measures how much the opinion profile x is exposed to negative-edge disagreement: it is large when enemies hold different opinions. The numerator $x^\top(\alpha L_{G^+} + D)x$ measures the stabilizing force acting against that disagreement: positive ties penalize disagreement among friends, while D anchors agents to external information. Thus $\beta_{\text{thresh}}(c)$ is the smallest ratio of stabilizing force to antagonistic exposure across all nontrivial directions. Stability holds exactly when the antagonism strength β is below this worst-case ratio.

This is the main constant-information result. It says that external information does not merely shift the eventual fixed point; it changes the range of antagonism levels for which a stable fixed point exists.

Corollary 5.4 (Media weight raises the stability threshold)**.** *If $c' \geq c$ coordinatewise, then $\beta_{\text{thresh}}(c') \geq \beta_{\text{thresh}}(c)$. In particular, increasing any media weight c_i weakly raises the stability threshold.*

Proof. Let $D' = \operatorname{diag}(c')$ and $A' = \alpha L_{G+} + D'$. If $c' \geq c$, then for every x,

$$x^\top A' x = x^\top A x + \sum_i (c'_i - c_i) x_i^2 \geq x^\top A x.$$

The denominator $x^\top L_- x$ is unchanged. Therefore the infimum defining $\beta_{\text{thresh}}(c')$ is weakly larger than the one defining $\beta_{\text{thresh}}(c)$. □

This corollary formalizes the anchoring role of information. An agent who places more weight on external information places less effective weight on the social feedback loop. That weakens the ability of negative edges to generate runaway disagreement. In this sense, media weight acts as a stabilizer.

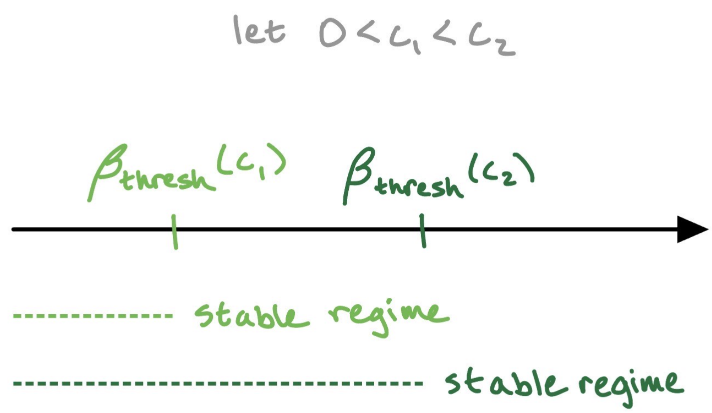

Figure 6: Media weight acts as a stabilizer.

Remark 5.5. The threshold $\beta_{\text{thresh}}(c)$ generalizes the no-information threshold β^* of Shi et al. (2019): setting $c = 0$ recovers the Shi threshold, up to differences in model parameterization. The point of the present formulation is that the stabilizing term now has two sources, αL_{G+} and D, so the model separates stabilization through positive social ties from stabilization through external information.

Remark 5.6 (Interpretation of Long-Run Fixed Point). The stability theorem identifies the fixed point in the stable regime:

$$x^* = (A - \beta L_-)^{-1} D\theta.$$

This formula describes how external information propagates through the signed network after all transient dynamics have disappeared.

Define

$$\tilde{B}(c,\beta) := (A - \beta L_-)^{-1}.$$

Then

$$x_i^* = \sum_j \tilde{B}(c,\beta)_{ij}\, c_j\theta_j.$$

Thus agent i's long-run opinion is a weighted combination of all agents' external information signals. The weights are not determined only by distance or friendship. They are entries of a signed Leontief-type inverse that incorporates attraction through αL_{G^+}, repulsion through βL_-, and anchoring through D. This is the constant-information benchmark for the time-varying analysis that follows.

6 Time-Varying Information and Frequency Response

The constant-information model analyzed in Section 5 leaves out an important empirical feature of many informational environments, namely, that signals change over time. For instance, news cycles may intensify and fade, partisan narratives might alternate, and financial markets might exhibit cyclical behavior. In this section, we propose a first extension of the repelling-with-information model to time-varying signals. Once signals vary over time, it becomes natural to study the entire opinion path rather than only a fixed point, and new questions (eg. about resonance or dynamic notions of centrality) can come into view.

Throughout this section, we write the forward recursion with time-varying signals as

$$x(t+1) = M'(\beta)x(t) + D\theta(t).$$

We focus in the main text on the forward causal model in the stable regime, because this is the setting used for the frequency-response and phase-lag optimization results below. A more general two-sided sequence-space formulation is possible: instead of specifying an initial condition and propagating forward, one can ask for a bounded bi-infinite opinion history consistent with the update rule at every date. This may admit richer or more elegant models, but a full exploration is beyond the scope of this thesis. We record this formulation in Appendix B.

6.1 Periodic Signals

The constant-information models discussed above are tractable and simple to understand, but they suffer from a certain lack of realism. Because their external information is constant, those models can, in essence, ask only a binary question: do opinions converge to a static fixed point, or do they become unstable?

In reality, however, we rarely observe systems that simply converge indefinitely towards some single fixed point (or at least, the long run behavior of such systems is rarely particularly interesting).

In order to address this, we study periodic signals. Periodic signals are the simplest time-varying inputs that are analytically tractable and not reducible to the constant-information case. They also have a natural interpretation.

For example, we might consider information about the stock market. Such information does not arrive once and then stay fixed: news, sentiment, and attention rise and fall over time. Prior work has shown that media pessimism is linked to stock-market activity (Tetlock, 2007), and that collective mood measured from Twitter is correlated with stock-market movements (Bollen et al., 2011). A trader's beliefs may therefore keep moving in a regular way even after the initial effect of her starting opinion has disappeared. Can we characterize how her opinion moves over time?

Social media platforms also have natural analogs here. Online news and social-media attention exhibit temporal rhythms and bursty cycles: memes and phrases rise and decay over the news cycle (Leskovec et al., 2009), and Twitter hashtags exhibit distinct dynamical classes of collective attention (Lehmann et al., 2012). "Hype cycles" on Twitter emerge and fade with some predictable frequency, invariably stirring up controversy on a semi-regular schedule (a schedule that, for better or for worse, pundits, podcasters, and politicians alike have – at least anecdotally – become rather adept at manipulating). How do these waves of periodic information affect the sentiment of viewers and voters? How do the behaviors induced by Twitter's high-frequency cycles differ from the behavior induced by its more "geriatric" counterparts (like Facebook, perhaps)?

In general if the external information itself keeps oscillating, how does the opinion vector behave? Does the opinion vector settle into a regular oscillation? Which parts of the network amplify that oscillation? Along which directions does the signal get amplified?

We emphasize that these questions are uniquely interesting to ask in antagonistic models – and in particular, the repelling model – because the antagonistic edges amplify information (and thus contribute to signal interference, constructive or otherwise) in ways that are potentially more complex than the simple weighted averaging native to nonnegative networks. It will also therefore be interesting to ask: what role does antagonism play in the behavior of the network/the amplification of the signal?

To begin answering these questions, we study the foundational case of single-frequency signals. Fix frequency $\omega \in [0, \pi]$ and let $v \in \mathbb{C}^n$ be a complex amplitude vector. The real signal schedule is

$$\theta(t) = \Re\big(ve^{i\omega t}\big). \tag{7}$$

We note that the complex notation serves only to drive the algebraic mechanics on the "backend"; the realized signal on the "frontend" is real. Its advantage is that time shifts act on $e^{i\omega t}$ by scalar multiplication, so the algebra is cleaner.

6.2 Periodic Steady States

We first ask whether a stable signed network has a well-defined long-run response to a single oscillating input. In the stock-market example, this asks whether traders' opinions eventually settle into a regular cycle induced by the news cycle, rather than converging to a fixed opinion vector. The answer is yes: after transients die out, the opinions oscillate at the same frequency as the signal.

Proposition 6.1 (Bounded steady-state response to a single oscillating signal)**.** *Assume $\rho(M'(\beta)) < 1$, and consider*

$$x(t+1) = M'(\beta)x(t) + D\,\Re\big(ve^{i\omega t}\big), \qquad t \geq 0. \tag{8}$$

Then there is a unique single-frequency periodic response of the form

$$x^{\text{per}}(t) = \Re\big(H_\beta(\omega)v\, e^{i\omega(t-1)}\big),$$

where

$$H_\beta(\omega) := (I - e^{-i\omega} M'(\beta))^{-1} D. \tag{9}$$

Moreover, every solution of (8) *converges to* $x^{\mathrm{per}}(t)$ *as* $t \to \infty$.

Proof. Consider the auxiliary complex system

$$z(t+1) = M'(\beta) z(t) + D v e^{i\omega t}.$$

Since $M'(\beta)$ and D are real, the real part of any solution is a solution of (8). We look for a same-frequency particular solution $z^{\mathrm{per}}(t) = h e^{i\omega t}$. Substitution gives

$$h e^{i\omega(t+1)} = M'(\beta) h e^{i\omega t} + D v e^{i\omega t},$$

hence

$$(e^{i\omega} I - M'(\beta)) h = Dv.$$

Thus

$$h = (e^{i\omega} I - M'(\beta))^{-1} D v = e^{-i\omega} (I - e^{-i\omega} M'(\beta))^{-1} D v = e^{-i\omega} H_\beta(\omega) v.$$

Taking real parts gives

$$x^{\mathrm{per}}(t) = \Re\big(H_\beta(\omega) v \, e^{i\omega(t-1)}\big).$$

It remains to show that all trajectories converge to this periodic response. Let

$$y(t) = x(t) - x^{\mathrm{per}}(t).$$

Since both terms satisfy the same input-driven recursion, their difference satisfies the homogeneous recursion

$$y(t+1) = M'(\beta) y(t).$$

Therefore $y(t) = M'(\beta)^t y(0)$. Since $\rho(M'(\beta)) < 1$, we have $M'(\beta)^t \to 0$, and hence $y(t) \to 0$. This proves convergence.

Finally, if $\tilde{x}^{\mathrm{per}}(t)$ is another single-frequency periodic response, then $x^{\mathrm{per}}(t) - \tilde{x}^{\mathrm{per}}(t)$ is both periodic and converges to zero. Hence it is identically zero. □

The proposition gives the object we will study below. The matrix $H_\beta(\omega)$ maps the input amplitude to the response amplitude with respect to the shifted carrier $e^{i\omega(t-1)}$. Equivalently, the phasor relative to $e^{i\omega t}$ is obtained from $B_\beta(\omega) v$, where

$$B_\beta(\omega) := (e^{i\omega} I - M'(\beta))^{-1} D = e^{-i\omega} H_\beta(\omega).$$

The two conventions differ only by a global phase factor, so all gain magnitudes below are unchanged.

6.3 Gain Along Eigenmodes and Arbitrary Opinion Profiles

We now ask a more fine-grained question. Once we know that the long-run periodic response exists, along which directions is the network most sensitive to an input signal of frequency ω? In a signed network, a direction may represent a consensus pattern, a disagreement pattern, or a more localized opinion profile. We therefore decompose the response into eigenmodes of the signed update matrix.

Definition 6.2 (Eigenmode)**.** Suppose $M'(\beta)$ is symmetric with orthonormal eigenpairs $(\lambda_k(\beta), q_k)_{k=1}^n$. The ***k*th eigenmode** is the one-dimensional direction $\operatorname{span}\{q_k\}$.

An eigenmode is a direction that the social update matrix scales without rotating into other directions. This makes eigenmodes the right coordinates for studying amplification: if an input signal loads onto one eigenmode, the network response along that mode can be described by a scalar.

Definition 6.3 (Frequency-specific gain of an eigenmode)**.** Fix an eigenmode q_k. Suppose the weighted input Dv has modal amplitude

$$b_k := q_k^\top Dv.$$

If the periodic steady state is written using the shifted carrier from Theorem 6.1, then its complex amplitude is $H_\beta(\omega)v$, and its modal amplitude along q_k is

$$a_k := q_k^\top H_\beta(\omega)v.$$

Thus $q_k^\top x^{\mathrm{per}}(t) = \Re(a_k e^{i\omega(t-1)})$, so $|a_k|$ is the amplitude of the opinion response along the kth eigenmode under this phase convention. The quantity $|b_k|$ is the amplitude of the weighted input signal along that same mode.

When $b_k \neq 0$, the frequency-ω gain of the kth eigenmode is the ratio

$$\frac{|a_k|}{|b_k|}.$$

The next proposition shows that this ratio is

$$g_k(\omega, \beta) := \frac{1}{|1 - e^{-i\omega}\lambda_k(\beta)|} = \frac{1}{\sqrt{1 + \lambda_k(\beta)^2 - 2\lambda_k(\beta)\cos\omega}}.$$

This is the steady-state amplitude amplification factor for information signals in the kth eigenmode.

Intuitively, "gain" compares two quantities measured in the same eigenmode: the amplitude of the input signal entering the system and the amplitude of the opinion response that remains in the periodic steady state.

Proposition 6.4 (Modal decomposition and frequency-ω gain)**.** *Assume $M'(\beta)$ is symmetric, with orthonormal eigenpairs $(\lambda_k(\beta), q_k)_{k=1}^n$. Then for every $v \in \mathbb{C}^n$,*

$$H_\beta(\omega)v = \sum_{k=1}^n \frac{q_k^\top Dv}{1 - e^{-i\omega}\lambda_k(\beta)}\, q_k.$$

Consequently, if $b_k = q_k^\top Dv$, then the response amplitude along q_k is

$$|a_k| = g_k(\omega, \beta)|b_k|,$$

where

$$g_k(\omega,\beta) = \frac{1}{|1-e^{-i\omega}\lambda_k(\beta)|} = \frac{1}{\sqrt{1+\lambda_k(\beta)^2 - 2\lambda_k(\beta)\cos\omega}}.$$

Proof. Let $y = H_\beta(\omega)v$. By definition,

$$(I - e^{-i\omega}M'(\beta))y = Dv.$$

Expand both sides in the orthonormal eigenbasis:

$$y = \sum_{k=1}^{n} a_k q_k, \qquad Dv = \sum_{k=1}^{n} (q_k^\top Dv) q_k.$$

Since $M'(\beta)q_k = \lambda_k(\beta)q_k$,

$$(I - e^{-i\omega}M'(\beta))q_k = (1 - e^{-i\omega}\lambda_k(\beta))q_k.$$

Matching coefficients gives

$$a_k(1 - e^{-i\omega}\lambda_k(\beta)) = q_k^\top Dv,$$

so

$$a_k = \frac{q_k^\top Dv}{1 - e^{-i\omega}\lambda_k(\beta)}.$$

This proves the modal decomposition. Taking absolute values gives the gain formula. Finally,

$$|1 - e^{-i\omega}\lambda_k(\beta)|^2 = 1 + \lambda_k(\beta)^2 - 2\lambda_k(\beta)\cos\omega,$$

which gives the stated expression. □

The modal formula answers a precise question: if the input loads onto a given eigenmode, how much is that modal coordinate amplified in the long-run periodic response? This is already a new kind of question relative to constant-information models, where the long-run object is a static fixed point rather than a periodic response with frequency-dependent amplitude.

Sometimes we do not care about a single eigenmode. For example, we may want to know how strongly the population responds along a chosen disagreement vector q, even if q is not exactly an eigenvector. The next corollary gives this more general directional response.

Corollary 6.5 (Gain along an arbitrary direction)**.** *Let $q \in \mathbb{R}^n$ be any unit vector. For input amplitude v, the shifted-carrier complex amplitude of the steady-state response along q is*

$$q^\top H_\beta(\omega)v = \sum_{k=1}^{n} \frac{(q^\top q_k)(q_k^\top Dv)}{1 - e^{-i\omega}\lambda_k(\beta)}.$$

Thus the response along an arbitrary direction is a superposition of the modal responses, weighted by the alignment of q with each eigenmode and by the loading of Dv onto that eigenmode.

Moreover, if one normalizes the input by $\|v\|_2 \le 1$, then the maximum possible response amplitude along q is

$$\max_{\|v\|_2 \le 1} |q^\top H_\beta(\omega)v| = \|H_\beta(\omega)^\top q\|_2.$$

Proof. The expansion follows immediately from Proposition 6.4 after applying $q^\top$. For the optimization statement, write

$$q^\top H_\beta(\omega)v = (H_\beta(\omega)^\top q)^\top v.$$

By Cauchy–Schwarz,

$$|(H_\beta(\omega)^\top q)^\top v| \le \|H_\beta(\omega)^\top q\|_2 \|v\|_2.$$

Equality is attained by choosing v proportional to $\overline{H_\beta(\omega)^\top q}$, when this vector is nonzero. □

This corollary has a direct interpretation. Fix a frequency ω, such as a regular cycle of market news. Fix a direction q, such as a disagreement vector that contrasts one group of agents with the rest of the population. Then $q^\top H_\beta(\omega)v$ is the shifted-carrier complex amplitude of the long-run opinion response along that disagreement direction. The formula above says that this response is built from the network's eigenmodes, with each mode weighted both by how much the input excites it and by how much it aligns with q.

6.4 Effect of Network Properties and Antagonism on Gain

The next corollary isolates the scalar part of the gain/"amplification" effect. It asks: for a fixed frequency, which eigenvalues produce the largest modal amplification?

Corollary 6.6 (The gain is maximized near $\cos\omega$). *Fix $\omega \in [0, \pi]$. As a function of a real scalar λ,*

$$g(\lambda, \omega) := \frac{1}{\sqrt{1 + \lambda^2 - 2\lambda\cos\omega}}$$

is maximized when λ is as close as possible to $\cos\omega$. Equivalently, among the eigenmodes of $M'(\beta)$, the modes whose eigenvalues are closest to $\cos\omega$ have the largest frequency-ω gain.

Proof. Maximizing $g(\lambda, \omega)$ is equivalent to minimizing

$$1 + \lambda^2 - 2\lambda\cos\omega = (\lambda - \cos\omega)^2 + \sin^2\omega.$$

The second term is independent of λ, so the expression is minimized when λ is closest to $\cos\omega$. □

Thus the network does not amplify all time-varying information in the same way. A slowly varying signal, with ω near 0, is most amplified by modes with eigenvalues near 1. A rapidly alternating signal, with ω near π, is most amplified by modes with eigenvalues near -1. Intermediate frequencies select intermediate parts of the spectrum.

The preceding calculation is a linear-systems statement, but the signed-network content enters through how antagonism moves the spectrum. In the repelling model, increasing β adds the positive semidefinite term βL_-, so modes exposed to negative edges move to the right in the spectrum.

Corollary 6.7 (Antagonism shifts modal gain through L_-). *Assume $\lambda_k(\beta)$ is a simple eigenvalue of $M'(\beta)$, with normalized eigenvector $q_k(\beta)$. Then*

$$\frac{d}{d\beta}\lambda_k(\beta) = q_k(\beta)^\top L_- q_k(\beta) \ge 0.$$

Consequently, increasing antagonism changes frequency-ω gain through the signed-network exposure term $q_k(\beta)^\top L_- q_k(\beta)$.

Proof. For a simple eigenvalue of the differentiable symmetric matrix $M'(\beta) = M'(0) + \beta L_-$, standard eigenvalue perturbation gives

$$\frac{d}{d\beta}\lambda_k(\beta) = q_k(\beta)^\top \frac{dM'(\beta)}{d\beta} q_k(\beta) = q_k(\beta)^\top L_- q_k(\beta).$$

Since $L_- \succeq 0$, this derivative is nonnegative. □

The endpoint frequencies $\omega = 0$ and $\omega = \pi$ correspond to constant and alternating signals, respectively. These endpoints make the effect of antagonism especially transparent.

Corollary 6.8 (Antagonism, constant signals, and alternating signals)**.** *Assume $\beta < \beta_{\text{thresh}}(c)$, so that $\rho(M'(\beta)) < 1$. For each fixed mode k,*

$$g_k(0,\beta) = \frac{1}{1-\lambda_k(\beta)}, \qquad g_k(\pi,\beta) = \frac{1}{1+\lambda_k(\beta)}.$$

Since the eigenvalues are nondecreasing in β by Theorem 5.2, the constant-signal gain $g_k(0,\beta)$ is nondecreasing in β, while the alternating-signal gain $g_k(\pi,\beta)$ is nonincreasing in β.

Proof. Substitute $\omega = 0$ and $\omega = \pi$ into the gain formula from Theorem 6.4. In the stable regime, $\lambda_k(\beta) \in (-1, 1)$, so both denominators are positive. As β increases, $\lambda_k(\beta)$ increases; hence $1-\lambda_k(\beta)$ decreases and $1 + \lambda_k(\beta)$ increases. □

	Constant signals ($\omega = 0$)	**Alternating signals** ($\omega = \pi$)
Low antagonism (β small)	Weaker amplification	Stronger amplification
High antagonism (β near $\beta_{\text{thresh}}(c)$)	Stronger amplification	Weaker amplification

The table should be read mode by mode. The total response also depends on how the weighted input Dv loads onto each eigenvector, but the effect of increasing antagonism on endpoint modal gain is monotone.

7 Optimizing Phase-Lagged Signals

The preceding subsection fixed a signal and asked how the network amplifies it. We now reverse the question. Suppose the network and frequency are fixed, but different agents receive the same information cycle with different magnitudes and lags. Which phase-lagged signal produces the largest steady-state response?

This question is natural in the oscillating-information setting. In a news-cycle example, the underlying story may be common, but agents may encounter it through different sources, at different times, and with different levels of attention. We model this by writing the complex signal amplitude as

$$a_i = r_i e^{i\phi_i}, \qquad r_i \geq 0, \qquad \phi_i \in [0, 2\pi).$$

Thus r_i is the magnitude of the signal received by agent i, while ϕ_i is the phase lag.

For a fixed frequency ω, the steady-state complex opinion amplitude is

$$z = (e^{i\omega} I - M'(\beta))^{-1} D a.$$

Equivalently, writing

$$B_\omega := (e^{i\omega} I - M'(\beta))^{-1} D,$$

we have $z = B_\omega a$. The real steady-state trajectory is

$$x_t^{\mathrm{ss}} = \mathrm{Re}(z e^{i\omega t}).$$

This section is exploratory. In full generality, optimizing over a is a nonconvex signal-design problem, and we do not attempt to solve it completely. Instead, we use it in two ways: first, we solve one small frustrated example analytically; second, we use numerical optimization to test whether the same kind of phase behavior appears in less symmetric networks.

The objective we focus on is time-averaged variance. Let

$$P = I - \frac{1}{n} \mathbf{1}\mathbf{1}^\top$$

be the projection away from consensus. We define

$$\overline{V}_{\mathrm{var}}(a) = \frac{1}{2\pi} \int_0^{2\pi} \frac{1}{n} \left\| P \, \mathrm{Re}(B_\omega a e^{i\tau}) \right\|_2^2 \, d\tau.$$

In the following numerical experiments, we maximize this objective under the constraint $\|r\|_\infty \leq 1$. This means that each agent may receive a signal of magnitude at most one.

7.1 A Solved Analytical Example: The Frustrated Triangle

We first consider the smallest signed graph in which pairwise antagonism cannot be resolved by a binary split: the complete negative triangle. This example is small enough to solve exactly, but already exhibits the main phenomenon we want to understand: the optimal phase lags need not be 0 and π.

In this example,

$$A^+ = 0, \qquad A^- = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix},$$

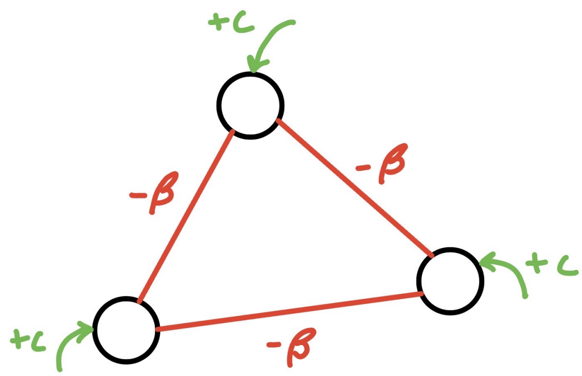

Figure 7: The frustrated triangle: three agents, with every edge negative.

so

$$L_+ = 0, \qquad L_- = \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix}.$$

Let $D = cI$, where $c > 0$, and set

$$W = I - \alpha L_+ + \beta L_- - D = (1-c)I + \beta L_-.$$

Assume $0 \le \beta < c/3$, so that the dynamics are stable.

Fix a frequency ω, and let the complex information signal be

$$\theta_t = a e^{i\omega t}, \qquad a_i = r_i e^{i\phi_i}, \qquad 0 \le r_i \le 1.$$

Let

$$z = (e^{i\omega} I - W)^{-1} D a$$

be the complex steady-state amplitude, so that

$$x_t^{\mathrm{ss}} = \mathrm{Re}(z e^{i\omega t}).$$

Finally, let

$$P = I - \frac{1}{3}\mathbf{1}\mathbf{1}^\top$$

be the projection onto the disagreement subspace, and define the phase-averaged variance objective

$$\overline{V}(a) = \frac{1}{2\pi}\int_0^{2\pi} \frac{1}{3}\left\| P\,\mathrm{Re}(z e^{i\tau})\right\|_2^2 \, d\tau.$$

Proposition 7.1 (A frustrated triangle admits a three-phase optimum)**.** *In the complete negative triangle setup above, $\overline{V}(a)$ is maximized by full magnitudes $r_1 = r_2 = r_3 = 1$ and phases separated by $120°$. Equivalently, the maximizers are exactly*

$$(a_1, a_2, a_3) = e^{i\psi}\left(1, e^{2\pi i/3}, e^{4\pi i/3}\right),$$

up to permutation, where $\psi \in [0, 2\pi)$ is arbitrary. In particular, the optimal lags are genuinely *three-phase and strictly dominate every binary-lag solution with phases restricted to* $\{0, \pi\}$.

Proof. Since the negative graph is the complete graph on three vertices,

$$L_- = 3P.$$

Therefore

$$W = (1-c)I + \beta L_- = (1-c)(I-P) + (1-c+3\beta)P.$$

Thus W has consensus eigenvalue

$$\lambda_{\text{cons}} = 1 - c$$

and disagreement eigenvalue

$$\lambda_{\text{dis}} = 1 - c + 3\beta.$$

The assumption $\beta < c/3$ gives $\lambda_{\text{dis}} < 1$, so the relevant steady-state response is well-defined.

Because P projects onto the disagreement subspace and W acts as scalar multiplication by λ_{dis} on this subspace,

$$Pz = P(e^{i\omega}I - W)^{-1}cIa = \frac{c}{e^{i\omega} - \lambda_{\text{dis}}}Pa.$$

Let

$$\gamma_\omega = \frac{c}{e^{i\omega} - \lambda_{\text{dis}}}.$$

Then

$$Pz = \gamma_\omega Pa.$$

Now compute the phase-averaged variance:

$$\overline{V}(a) = \frac{1}{2\pi}\int_0^{2\pi} \frac{1}{3}\left\|\operatorname{Re}(Pze^{i\tau})\right\|_2^2 \, d\tau.$$

Using the identity

$$\frac{1}{2\pi}\int_0^{2\pi} \left\|\operatorname{Re}(ue^{i\tau})\right\|_2^2 \, d\tau = \frac{1}{2}\|u\|_2^2$$

for any $u \in \mathbb{C}^3$, we get

$$\overline{V}(a) = \frac{1}{6}\|Pz\|_2^2 = \frac{|\gamma_\omega|^2}{6}\|Pa\|_2^2.$$

Thus maximizing $\overline{V}(a)$ is equivalent to maximizing

$$\|Pa\|_2^2$$

subject to $|a_i| \le 1$.

Since

$$P = I - \frac{1}{3}\mathbf{1}\mathbf{1}^\top,$$

we have

$$\|Pa\|_2^2 = \|a\|_2^2 - \frac{1}{3}|\mathbf{1}^\top a|^2 = \sum_{i=1}^{3}|a_i|^2 - \frac{1}{3}|a_1 + a_2 + a_3|^2.$$

Because $|a_i| \le 1$, the first term is at most 3, and the second term is nonpositive. Hence

$$\|Pa\|_2^2 \le 3.$$

Equality holds if and only if

$$|a_1| = |a_2| = |a_3| = 1$$

and

$$a_1 + a_2 + a_3 = 0.$$

Three unit complex numbers sum to zero exactly when they are separated by angles of $2\pi/3$. Hence the maximizers are

$$(a_1, a_2, a_3) = e^{i\psi}\left(1, e^{2\pi i/3}, e^{4\pi i/3}\right),$$

up to permutation.

It remains to compare against binary-lag solutions. If $a_i \in [-1, 1]$, then the maximum of the convex quadratic $a^\top Pa$ over the cube $[-1, 1]^3$ is attained at a vertex. The best binary vertices are permutations of

$$a = (1, 1, -1).$$

For this vector,

$$\|Pa\|_2^2 = \|a\|_2^2 - \frac{1}{3}(\mathbf{1}^\top a)^2 = 3 - \frac{1}{3} = \frac{8}{3}.$$

The three-phase solution gives

$$\|Pa\|_2^2 = 3.$$

Therefore

$$3 > \frac{8}{3},$$

so the three-phase solution strictly dominates every binary-lag solution. □

This example is deliberately small. Its value is that the conclusion can be seen exactly. A structurally balanced two-camp network often suggests a binary phase split: one camp at phase 0, the other at phase π. The negative triangle cannot satisfy all pairwise antagonistic relations with only two phases. The variance-maximizing signal instead places the three agents at evenly spaced points on the unit circle.

One might worry that this example is too symmetric. The proof uses the identity $L_- = 3P$, which makes the algebra collapse to a geometric problem about three complex numbers. Perhaps surprisingly, however, the numerical experiments below provide some compelling initial evidence that similar qualitative behavior persists even in larger signed networks where no such exact simplification is available.

7.2 Numerical Methodology

All numerical experiments in this section use the Shi repelling-with-information form

$$M_\beta = I - \alpha L_+ + \beta L_- - D,$$

where L_+ and L_- are constructed from the positive and negative edges of the signed graph. The code used for these experiments is available at https://github.com/sophiapi-nu26/oscillatory-optimizer,

and an interactive version of the optimizer is available at

https://spi-net.streamlit.app/.

The numerical pipeline has four steps. First, we generate a signed graph and construct M_β. Second, we check that the selected parameters are in the stable regime by computing $\rho(M_\beta)$. If $\rho(M_\beta) \geq 1$, the steady-state response need not exist, so the experiment is not run. Third, for a chosen frequency ω, we compute the frequency-response matrix

$$B_\omega = (e^{i\omega} I - M_\beta)^{-1} D.$$

Fourth, we optimize over phase-lagged information signals of the form

$$a_i = r_i e^{i\phi_i}, \qquad r_i \geq 0, \qquad \phi_i \in [0, 2\pi).$$

Although the code supports several objectives, including total norm and projection onto a prescribed direction, the experiments reported here focus on the time-averaged variance objective. Let

$$P = I - \frac{1}{n}\mathbf{1}\mathbf{1}^\top.$$

The real steady-state trajectory induced by a is

$$x_t^{\mathrm{ss}} = \mathrm{Re}(B_\omega a e^{i\omega t}).$$

The objective is

$$\overline{V}_{\mathrm{var}}(a) = \frac{1}{2\pi}\int_0^{2\pi} \frac{1}{n}\left\| P\,\mathrm{Re}(B_\omega a e^{i\tau})\right\|_2^2 \, d\tau.$$

Using the periodic time-average identities

$$\overline{\cos^2} = \overline{\sin^2} = \frac{1}{2}, \qquad \overline{\cos\sin} = 0,$$

this objective reduces to the Hermitian quadratic form

$$\overline{V}_{\mathrm{var}}(a) = a^* Q a, \qquad Q = \frac{1}{2n} B_\omega^* P B_\omega.$$

Thus the optimization problem is

$$\max_{a\in\mathbb{C}^n} a^* Q a \qquad \text{subject to} \qquad a_i = r_i e^{i\phi_i}, \quad \|r\|_\infty \leq 1.$$

In the figures below, node size represents the optimized magnitude r_i, and node color or angular position represents the optimized phase ϕ_i.

Algorithm 1 Oscillatory information optimizer

Require: Signed adjacencies A^+, A^-, parameters α, β, c, ω, objective, norm constraint

Compute $L_+ = \mathrm{diag}(A^+\mathbf{1}) - A^+$ and $L_- = \mathrm{diag}(A^-\mathbf{1}) - A^-$

Set $D = \mathrm{diag}(c)$

Construct $M_\beta = I - \alpha L_+ + \beta L_- - D$

Compute diagnostics: $\rho(M_\beta)$, spectral edges, and $\mathrm{cond}(e^{i\omega} I - M_\beta)$

if $\rho(M_\beta) \geq 1$ **then**

  Stop; the steady-state response is not guaranteed to exist

end if

Compute $B_\omega = (e^{i\omega} I - M_\beta)^{-1} D$

Build the Hermitian objective matrix Q ▷ For variance, $Q = \frac{1}{2n} B_\omega^* P B_\omega$

Solve $\max a^* Q a$ subject to $a_i = r_i e^{i\phi_i}$ and the chosen norm constraint on r

Recover $r_i = |a_i|$ and $\phi_i = \arg(a_i)$

Compute $x_t^{\mathrm{ss}} = \mathrm{Re}(B_\omega a e^{i\omega t})$

Plot the signed network, optimized magnitudes, optimized phases, and objective value over time

For the ℓ_∞ constraint used in the main experiments, the optimizer uses a phase-projected power iteration with multiple random restarts. Starting from random phases on the complex unit circle, each iteration computes Qa and updates each coordinate of a to have the phase of the corresponding coordinate of Qa, while keeping $|a_i| \leq 1$. Multiple restarts are important because binary real configurations can be fixed points of the iteration even when they are not globally optimal.

7.3 Numerical Experiments

We now report a few experiments that provide evidence for multi-phase lags and raise other interesting questions that we have yet to formalize into analytical theory. We include them because they show that the phase-lag problem has structure beyond toy examples like our frustrated triangle, and provide tangible examples upon which to build future conjectures.

Experiment	n	Edge description	α	β	c	ω
3-block SBM	15	friendly in-block only	0.05	0.01	0.15	0.62
All-antagonistic cycle	9	all edges negative	0.05	0.01	0.15	0.5
Lattice	16	rows/friend edges as specified	0.05	0.01	0.15	$0.48, 0.84, 1.32$

Table 3: Numerical experiment specifications. All experiments maximize time-averaged variance under $\|r\|_\infty \leq 1$.

7.3.1 Three-Block Signed SBM

Edges within a block are positive and edges across blocks are negative. The optimized lags form an approximate three-spoke pattern, with agents in the same block assigned similar phases. This is the closest stochastic analogue of the frustrated-triangle calculation: the three groups cannot all be pairwise opposite, so the optimizer separates them around the circle rather than collapsing to a binary split.

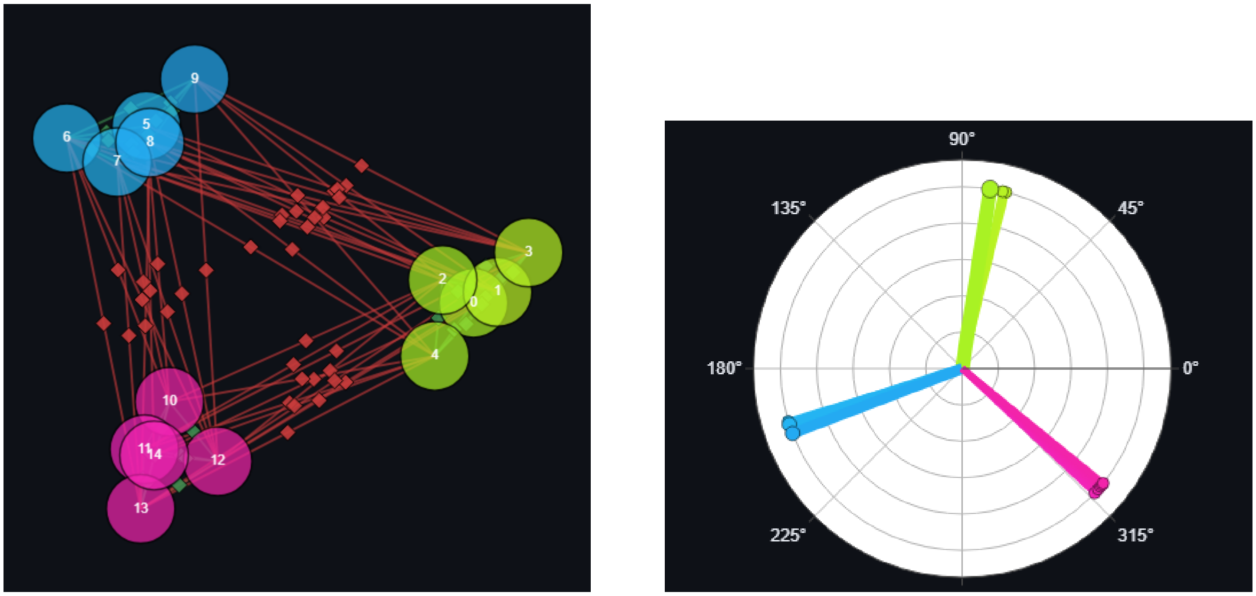

(a) Signed 3-community SBM.

(b) Optimized 3-community SBM lags.

Figure 8: Three-block signed SBM: variance maximization under $\|r\|_\infty \leq 1$ gives a three-spoke phase pattern.

7.3.2 All-Antagonistic k-Cycles

When every edge of a k-cycle is negative, the optimizer produces k spokes around the complex unit circle. This suggests that the negative-triangle calculation is not only an $n = 3$ artifact. Odd antagonistic cycles are the natural next test case because they generalize the frustration already present in the triangle.

7.3.3 Signed Lattices

In rectangular lattices with positive edges across rows and negative edges otherwise, low frequencies produce row-based spokes, while high frequencies produce column-based spokes. Adding positive edges between two rows draws the corresponding spokes closer together, consistent with positive ties acting as a phase-aligning force. This result is interesting for at least two reasons. First, it demonstrates that, perhaps counterintuitively, there are regimes in which the class of variance-maximizing phase lag assignments actually assigns heterogeneous phase lags within "friendly" groups instead of across "friendly" groups. Second, it demonstrates that the optimal phase lag assignment depends nontrivially on the frequency ω of the ground truth signal, which raises further questions about the interpretation of signal frequency and its relationship with phase lags.

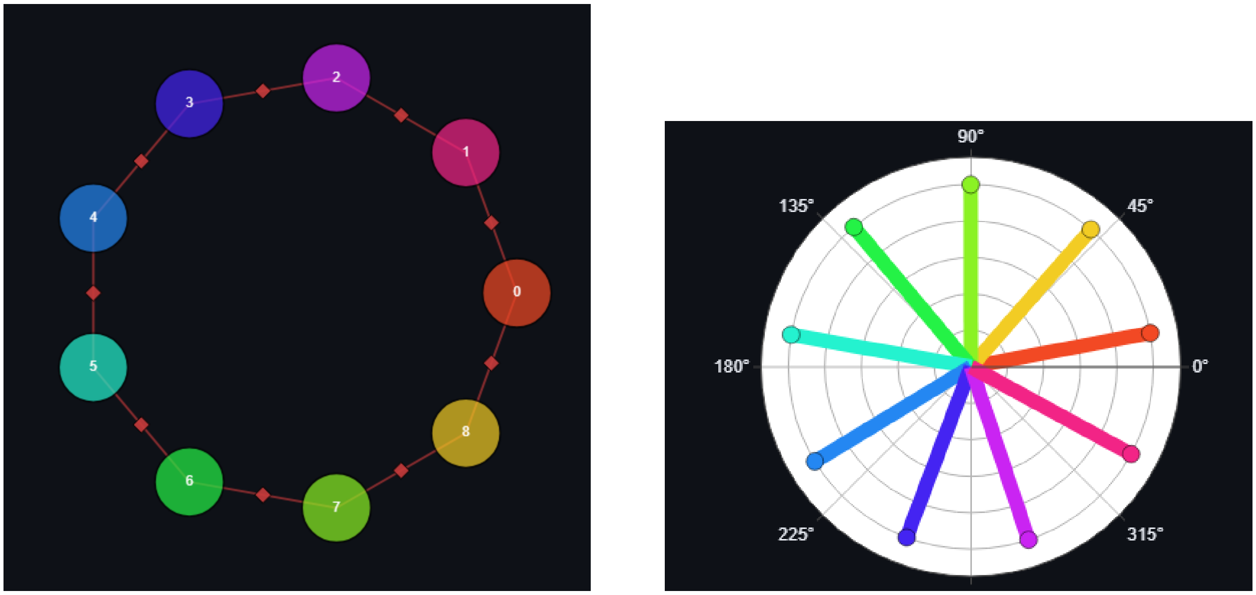

(a) Antagonistic 9-cycle.

(b) Optimized 9-cycle lags.

Figure 9: Antagonistic 9-cycle: variance maximization under $\|r\|_\infty \leq 1$ gives an (evenly spaced) 9-spoke phase pattern.

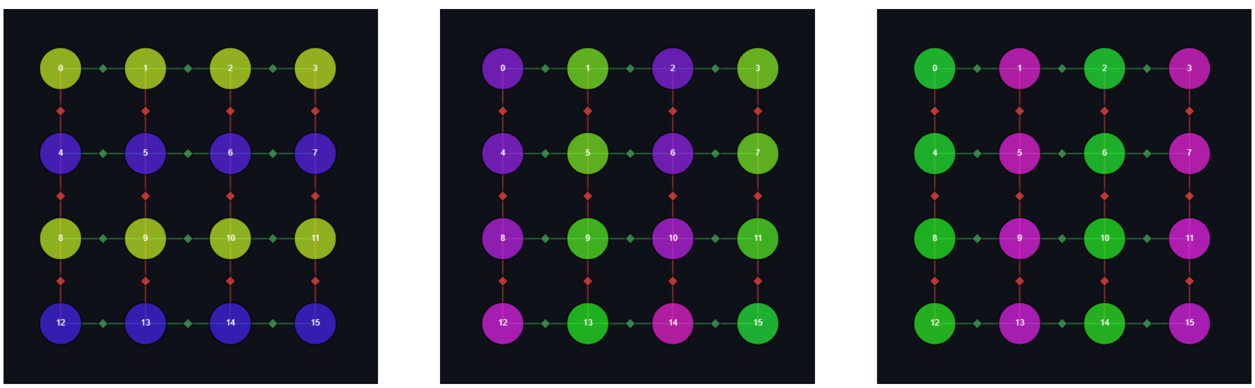

(a) Phase lag pattern on a lattice at frequency $\omega = 0.48$.

(b) Phase lag pattern on a lattice at frequency $\omega = 0.84$.

(c) Phase lag pattern on a lattice at frequency $\omega = 1.32$.

Figure 10: Lattice network: variance maximization under $\|r\|_\infty \leq 1$ give qualitatively different phase lag profiles at different frequencies.

8 Discussion and Open Questions

This thesis began with the simple observation that allowing negative social weights is not just a cosmetic modification of DeGroot learning. The opposing model shows one limitation of a first signed extension: without structural balance, opinions decay to zero, while under structural balance the dynamics reduce to DeGroot after a change of sign convention. The repelling model is different. Negative ties cannot be removed by relabeling; they create a destabilizing force that competes with positive ties and external information.

For constant information, this competition is captured by the threshold

$$\beta_{\text{thresh}}(c) = \inf_{x^\top L_- x > 0} \frac{x^\top(\alpha L_{G^+} + D)x}{x^\top L_- x}.$$

This formula gives a behavioral interpretation of stability: the numerator measures stabilizing forces from positive ties and external information, while the denominator measures exposure to antagonistic disagreement. The model is stable exactly when the strength of antagonism remains below the weakest stabilizing ratio.

The stability theorem is the main mathematical anchor of the thesis. The oscillating-information analysis is the main conceptual extension. Once the external signal varies over time, the long-run object is no longer a static fixed point but a periodic response. This makes it possible to ask frequency-specific questions: which modes respond to slow information cycles, which modes respond to alternating signals, and how does increasing antagonism change that response? The frequency-response matrix $H_\beta(\omega)$ and modal gain formula

$$g_k(\omega, \beta) = \frac{1}{|1 - e^{-i\omega}\lambda_k(\beta)|}$$

give a first answer. They give the stable repelling model a frequency-response interpretation: different eigenmodes are amplified at different frequencies, and the signed-network content enters through how antagonism shifts the spectrum.

The numerical experiments suggest that this frequency-response perspective is only the beginning. Once agents are allowed to receive the same oscillating information with different phase lags, the model produces phase patterns that do not appear in constant-information models. In the frustrated triangle, the optimal variance-maximizing signal has three phases separated by $120°$. The experiments on signed block models, antagonistic cycles, and lattices suggest that similar multi-spoke phase patterns persist in larger and less symmetric networks.

Several questions remain open.

1. **Frequency-specific centrality.** The matrix $H_\beta(\omega)$ suggests natural source and receiver centrality notions: some agents are influential because signals sent through them excite large network responses, while others are sensitive because their opinions respond strongly to signals elsewhere. A systematic theory of these frequency-specific centralities remains to be developed, especially their relationship to Katz–Bonacich centrality in the constant-information case.

2. **The role of antagonism in the optimizer.** The modal gain formula explains how β shifts eigenvalues and therefore changes amplification along fixed modes. The numerical optimizer, however, often changes the phase-lag pattern itself as β changes. Understanding when increasing antagonism creates new phase clusters, merges old ones, or changes the relevant partition of the network is still open.

3. **Beyond symmetric unweighted signed graphs.** The main results use symmetric Shi-type matrices with uniform positive and negative edge weights $+\alpha$ and $-\beta$. A more realistic model would allow directed influence, heterogeneous edge weights, and asymmetric information weights. In that setting, eigenmodes need not be orthogonal and modal gain may need to be replaced by singular-value or pseudospectral notions of response.

4. **Using the doubly infinite model.** The appendix formulates a two-sided sequence-space version of the model. In the stable regime, it agrees with the forward causal interpretation. Outside that regime, bounded histories may still exist but can depend on future signal values. This may be useful for modeling anticipatory or equilibrium notions of opinion formation, but the present thesis does not exploit that structure.

5. **Network optimization.** This thesis optimizes phase-lagged signals for a fixed signed network. A harder inverse question is to optimize the network itself: which signed topology, or which allocation of positive and negative weights, maximizes response along a prescribed disagreement profile? This would turn the frequency-response framework into a design problem rather than only an analysis tool.

This thesis has introduced a new regime of signed opinion dynamics in which antagonism, external information, and time variation interact. Opposing dynamics clarify why structural balance is such a strong assumption; repelling dynamics with information produce a sharp stability threshold; oscillating information turns the model into a frequency-response problem; and phase-lag optimization reveals new phenomena tied to signed frustration. Together, these results suggest that signed networks with dynamic external information form a richer class of models than the constant-information and unsigned settings capture.

References

Johan Bollen, Huina Mao, and Xiao-Jun Zeng. Twitter mood predicts the stock market. **Journal of Computational Science**, 2(1):1–8, 2011. doi: 10.1016/j.jocs.2010.12.007.

William J. Brady, Julian A. Wills, John T. Jost, Joshua A. Tucker, and Jay J. Van Bavel. Emotion shapes the diffusion of moralized content in social networks. **Proceedings of the National Academy of Sciences**, 114(28):7313–7318, 2017. doi: 10.1073/pnas.1618923114.

Dorwin Cartwright and Frank Harary. Structural balance: A generalization of heider's theory. **Psychological Review**, 63(5):277–293, 1956. doi: 10.1037/h0046049. URL https://doi.org/10.1037/h0046049.

Morris H. DeGroot. Reaching a consensus. **Journal of the American Statistical Association**, 69(345):118–121, 1974. doi: 10.1080/01621459.1974.10480137. URL https://doi.org/10.1080/01621459.1974.10480137.

Janette Lehmann, Bruno Gonçalves, José J. Ramasco, and Ciro Cattuto. Dynamical classes of collective attention in twitter. In **Proceedings of the 21st International Conference on World Wide Web**, pages 251–260, 2012. doi: 10.1145/2187836.2187871.

Sebastiano Della Lena, Luca Paolo Merlino, and Yves Zenou. Affective polarization, media outlets, and opinion dynamics. **SSRN Electronic Journal**, 2023. doi: 10.2139/ssrn.4595200. URL https://ssrn.com/abstract=4595200. Working paper.

Jure Leskovec, Lars Backstrom, and Jon Kleinberg. Meme-tracking and the dynamics of the news cycle. In **Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining**, pages 497–506, 2009. doi: 10.1145/1557019.1557077.

Guodong Shi, Claudio Altafini, and John S. Baras. Dynamics over signed networks. **SIAM Review**, 61(2):229–257, 2019. doi: 10.1137/17M1134172. URL https://doi.org/10.1137/17M1134172.

Paul C. Tetlock. Giving content to investor sentiment: The role of media in the stock market. **The Journal of Finance**, 62(3):1139–1168, 2007. doi: 10.1111/j.1540-6261.2007.01232.x.

Soroush Vosoughi, Deb Roy, and Sinan Aral. The spread of true and false news online. **Science**, 359(6380):1146–1151, 2018. doi: 10.1126/science.aap9559.

A Spectral Tools

This appendix collects several standard facts used in the main text. None of the results below are meant to be the main contribution of the thesis; they are included to make the spectral arguments self-contained. The organizing question is simple: when does a signed update matrix produce bounded, convergent, or divergent dynamics?

A.1 Laplacian Preliminaries

We first record the basic positivity property of graph Laplacians. This is used repeatedly when comparing the stabilizing contribution of positive ties to the destabilizing contribution of negative ties.

Lemma A.1 (Laplacians are PSD)**.** *Let G^+ and G^- be undirected graphs on n nodes. Then L_{G^+} and L_- are positive semidefinite.*

Proof. For any $x \in \mathbb{R}^n$,

$$x^\top L_{G^+} x = \sum_{\{i,j\}\in E^+} (x_i - x_j)^2 \geq 0.$$

This follows by expanding $x^\top (D_{G^+} - A_{G^+})x$ and using the symmetry of G^+. The same computation applies to L_- with E^- in place of E^+. □

Corollary A.2. *$L^r_{G^-} = -L_-$ is negative semidefinite.*

Proof. This follows immediately from Lemma A.1. □

A.2 Structural Balance and the Operator Norm

Structural balance is the condition that the signs in a signed network can be removed by flipping the signs of one group of agents. Equivalently, after a diagonal sign change, the signed matrix becomes entrywise nonnegative. We record only the direction of the operator-norm statement needed for intuition in the main text.

Lemma A.3 (Structural balance implies operator norm equality)**.** *Let $\tilde{W} \in \mathbb{R}^{n\times n}$, and let $W := |\tilde{W}|$ be its entrywise absolute value. If $\tilde{W}$ is structurally balanced, then*

$$\left\|\tilde{W}\right\|_{\mathrm{op}} = \|W\|_{\mathrm{op}}.$$

Proof. If $\tilde{W}$ is structurally balanced, then there is a diagonal sign matrix D_S, with diagonal entries in $\{\pm 1\}$, such that

$$D_S \tilde{W} D_S^{-1} = W \geq 0.$$

Equivalently, $\tilde{W} = D_S W D_S^{-1}$. Since D_S is orthogonal, multiplication by D_S does not change Euclidean norms. Therefore

$$\left\|\tilde{W}\right\|_{\mathrm{op}} = \left\|D_S W D_S^{-1}\right\|_{\mathrm{op}} = \|W\|_{\mathrm{op}}.$$

□

Remark A.4. The converse is false in general: equality of operator norms does not imply structural balance. For example, the all-negative triangle has the same operator norm as its entrywise absolute value, but its sign pattern is not structurally balanced. In the main text, structural balance is therefore used through the explicit diagonal sign-change condition, not through the converse of this lemma.

A.3 Boundedness and the Spectral Radius

The next facts are basic stability tests. They are useful because signed row sums alone do not control the spectral radius. What controls the spectral radius directly is an absolute row-sum bound.

Fact A.5 (Absolute row-stochasticity implies bounded powers)**.** *If*

$$|||W|||_\infty = \max_i \sum_j |w_{ij}| \le 1,$$

then $\|W^t\|_\infty \le 1$ *for all* $t \ge 0$*. In particular, the sequence* $(W^t)_{t\ge 0}$ *is bounded.*

Proof. By submultiplicativity of the ∞-norm,

$$\|W^t\|_\infty \le \|W\|_\infty^t \le |||W|||_\infty^t \le 1.$$

□

Corollary A.6. *Under the same condition, every eigenvalue of* W *satisfies* $|\lambda| \le 1$*. Moreover, any eigenvalue on the unit circle is semisimple.*

Proof. If W had an eigenvalue $|\lambda| > 1$, then $\|W^t\|$ would grow exponentially along the corresponding generalized eigenspace. If W had a nonsemisimple eigenvalue with $|\lambda| = 1$, then the Jordan block would produce polynomial growth in $\|W^t\|$. Both possibilities contradict boundedness of (W^t). □

A.4 Row-Sum and Spectral-Radius Regimes

The long-run behavior of a linear update $x(t+1) = Wx(t)$, or an affine update $x(t+1) = Wx(t)+u$, depends primarily on $\rho(W)$. The row-sum condition helps in nonnegative models, but it is less decisive for signed matrices.

There are three main spectral regimes.

- **$\rho(W) < 1$.** The homogeneous system contracts globally: $W^t \to 0$. The matrix $I - W$ is invertible, and with constant input u, the affine system converges to the unique fixed point

  $$x^* = (I - W)^{-1}u.$$

- **$\rho(W) = 1$.** Long-run behavior is governed by the unit-circle eigenspaces. If 1 is the only unit-circle eigenvalue and is semisimple, trajectories converge to the 1-eigenspace. In opinion models this may correspond to consensus or bipartite consensus. With constant inputs, fixed points may be nonunique or nonexistent, depending on the compatibility of u with the left 1-eigenspace.

- $\rho(W) > 1$**.** At least one mode expands. For generic initial conditions or inputs with nonzero projection onto an unstable eigendirection, the system diverges.

The important caution for signed networks is that a signed row-sum condition such as $W\mathbf{1} = c\mathbf{1}$ does not by itself control $\rho(W)$. A signed matrix can have small signed row sums while still having large absolute row sums and unstable eigenvalues. The reliable sufficient condition is instead

$$\||W|\|_\infty < 1 \qquad \Longrightarrow \qquad \rho(W) < 1.$$

Structural balance gives a second route back to the nonnegative case: if $\tilde{W}$ is structurally balanced, then it is similar by a diagonal sign change to $|\tilde{W}|$. In that case, spectral questions about $\tilde{W}$ reduce to spectral questions about a nonnegative matrix.

B Doubly Infinite Sequence-Space Model

The main text studies time-varying information in the forward causal model

$$x(t+1) = M'(\beta)x(t) + D\theta(t), \qquad t \geq 0. \tag{10}$$

This is the natural formulation when one specifies an initial opinion vector and asks what happens as time moves forward. In this appendix, however, we use the convention that $\theta(t)$ is the signal entering the update that produces $x(t)$. This differs from the forward convention by one time index, but the difference is not substantive for the solvability results below.

There is also a more symmetric way to write the same kind of dynamics. Instead of starting from one initial condition, we can package the entire opinion trajectory into a single bi-infinite object. This formulation is not needed for the main frequency-response calculations, but it clarifies what changes when we ask for bounded histories rather than forward trajectories.

Let

$$\mathcal{H}_\infty := \ell^\infty(\mathbb{Z};\mathbb{R}^n) = \left\{ X = (x(t))_{t\in\mathbb{Z}} : \sup_{t\in\mathbb{Z}} \|x(t)\| < \infty \right\}.$$

Define the right-shift operator $R : \mathcal{H}_\infty \to \mathcal{H}_\infty$ by

$$(RX)(t) := x(t-1).$$

For any matrix $B \in \mathbb{R}^{n\times n}$, let $\hat{B}$ denote the induced operator on histories:

$$(\hat{B}X)(t) := Bx(t).$$

Thus $\hat{B}$ applies B separately at each time. With this notation, the doubly infinite version of (10) is

$$X = \hat{M}'(\beta)RX + \hat{D}\Theta, \tag{11}$$

where $X = (x(t))_{t\in\mathbb{Z}}$ and $\Theta = (\theta(t))_{t\in\mathbb{Z}}$ are bounded histories.

The difference from the forward model is conceptual. In (10), one specifies $x(0)$ and propagates forward. In (11), there is no distinguished initial date; instead, one asks for a bounded bi-infinite history satisfying the update rule at every time. When all modes of $M'(\beta)$ are forward-stable, these two perspectives are compatible. When some modes are forward-unstable, bounded bi-infinite histories may still exist, but they generally depend on future signal values rather than only on past ones.

B.1 Reduction to the forward constant-signal model

The doubly infinite formulation induces the same recursion as the ordinary forward model on the nonnegative-time tail. To recover the constant-information dynamics from Section 5, use a step signal: no external signal before time 0, and a constant signal thereafter.

Proposition B.1 (Step histories induce the forward constant-signal recursion). *Fix $\theta \in \mathbb{R}^n$, and define the signal history $\Theta^\theta \in \mathcal{H}_\infty$ by*

$$\Theta^\theta(t) := \begin{cases} 0, & t < 0, \\ \theta, & t \geq 0. \end{cases}$$

Let $X = (x(t))_{t\in\mathbb{Z}} \in \mathcal{H}_\infty$ solve the history equation

$$X = \hat{M}'(\beta)RX + \hat{D}\Theta^\theta.$$

Then its nonnegative-time tail satisfies

$$x(t) = M'(\beta)x(t-1) + D\theta, \qquad t \geq 0.$$

Equivalently, if

$$y(t) := x(t-1), \qquad t \geq 0,$$

then

$$y(t+1) = M'(\beta)y(t) + D\theta, \qquad t \geq 0.$$

Thus the nonnegative-time tail of the doubly infinite model satisfies the same recursion as the forward causal constant-signal model, with initial state $y(0) = x(-1)$ determined by the chosen bounded prehistory. This should not be read as allowing an arbitrary forward initial condition inside the two-sided bounded-history formulation.

Proof. Write the history equation coordinatewise. Since

$$X = \hat{M}'(\beta)RX + \hat{D}\Theta^\theta,$$

its tth coordinate is

$$x(t) = M'(\beta)x(t-1) + D\Theta^\theta(t), \qquad t \in \mathbb{Z}.$$

For $t \geq 0$, $\Theta^\theta(t) = \theta$, so

$$x(t) = M'(\beta)x(t-1) + D\theta, \qquad t \geq 0.$$

Now define $y(t) := x(t-1)$ for $t \geq 0$. Then

$$y(t+1) = x(t) = M'(\beta)x(t-1) + D\theta = M'(\beta)y(t) + D\theta.$$

The initial condition is $y(0) = x(-1)$. □

Remark B.2. If instead $\Theta(t) \equiv \theta$ is a constant signal history and one looks for a constant solution $X(t) \equiv x^*$, then the history equation reduces to

$$x^* = M'(\beta)x^* + D\theta,$$

equivalently

$$(I - M'(\beta))x^* = D\theta.$$

Thus constant histories recover the fixed-point equation from Section 5, while step histories recover the forward causal recursion on the nonnegative-time tail, subject to the compatibility imposed by the bounded prehistory.

B.2 Existence of bounded solution histories

We now ask when the history equation (11) has a unique bounded solution for every bounded signal history. The scalar case contains the main idea.

Lemma B.3 (A scalar two-sided recurrence)**.** *Fix $\mu \in \mathbb{R}$, and consider the scalar equation*

$$z(t) - \mu z(t-1) = u(t), \qquad t \in \mathbb{Z}, \tag{12}$$

with $u \in \ell^\infty(\mathbb{Z};\mathbb{R})$.

1. *If $|\mu| < 1$, then for every bounded u there is a unique bounded solution:*

$$z(t) = \sum_{k=0}^{\infty} \mu^k u(t-k).$$

2. *If $|\mu| > 1$, then for every bounded u there is again a unique bounded solution:*

$$z(t) = -\sum_{k=1}^{\infty} \mu^{-k} u(t+k).$$

3. *If $|\mu| = 1$, then (12) is not well-posed on $\ell^\infty(\mathbb{Z};\mathbb{R})$: bounded solutions are either nonunique or fail to exist for some bounded inputs.*

Proof. First suppose $|\mu| < 1$. Define

$$z(t) := \sum_{k=0}^{\infty} \mu^k u(t-k).$$

Since u is bounded and $\sum_k |\mu|^k < \infty$, the series converges absolutely and uniformly in t, so z is bounded. An index shift gives

$$z(t) - \mu z(t-1) = \sum_{k=0}^{\infty} \mu^k u(t-k) - \sum_{k=1}^{\infty} \mu^k u(t-k) = u(t).$$

For uniqueness, suppose $y(t) - \mu y(t-1) = 0$. Then $y(t) = \mu^m y(t-m)$ for every $m \geq 0$. Since y is bounded and $|\mu| < 1$,

$$|y(t)| \leq |\mu|^m \|y\|_\infty \to 0,$$

so $y(t) = 0$.

Now suppose $|\mu| > 1$. Define

$$z(t) := -\sum_{k=1}^{\infty} \mu^{-k} u(t+k).$$

This series again converges absolutely and uniformly because $|\mu|^{-1} < 1$. A direct computation gives

$$\mu z(t-1) = -\sum_{k=1}^{\infty} \mu^{-k+1} u(t-1+k) = -\sum_{m=0}^{\infty} \mu^{-m} u(t+m),$$

and therefore

$$z(t) - \mu z(t-1) = -\sum_{k=1}^{\infty} \mu^{-k} u(t+k) + \sum_{m=0}^{\infty} \mu^{-m} u(t+m) = u(t).$$

The homogeneous equation is unique by iterating forward:

$$y(t) = \mu^{-m} y(t+m) \to 0.$$

Finally, if $|\mu| = 1$, well-posedness fails. For $\mu = 1$, every constant sequence solves the homogeneous equation, so uniqueness fails; also, $u(t) \equiv 1$ gives $z(t) - z(t-1) = 1$, whose solutions grow linearly. For $\mu = -1$, the homogeneous equation has the bounded nonzero solution $z(t) = (-1)^t$, and $u(t) = (-1)^t$ similarly produces no bounded solution after reducing to a linear-growth recursion. Hence the boundary case $|\mu| = 1$ is exactly where bounded-history well-posedness breaks down. □

The point of the scalar lemma is that bounded two-sided solvability is not the same as forward stability. If $|\mu| < 1$, the solution is causal and depends on past inputs. If $|\mu| > 1$, a bounded two-sided solution still exists, but it is anti-causal: it depends on future inputs. The boundary case $|\mu| = 1$ is the obstruction.

Theorem B.4 (History-space criterion for bounded solvability)**.** *Assume $M'(\beta)$ is symmetric. Then the following are equivalent:*

1. *For every bounded signal history $\Theta \in \mathcal{H}_\infty$, the history equation* (11) *has a unique bounded solution $X \in \mathcal{H}_\infty$.*
2. *$\sigma(M'(\beta)) \cap \{z \in \mathbb{C} : |z| = 1\} = \varnothing$.*
3. *Since $M'(\beta)$ is real symmetric, equivalently $\pm 1 \notin \sigma(M'(\beta))$.*

Proof. Diagonalize

$$M'(\beta) = Q\Lambda Q^\top, \qquad \Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_n),$$

with Q orthogonal. For a history $X = (x(t))_{t\in\mathbb{Z}}$, define the modal history $Z = (z(t))_{t\in\mathbb{Z}}$ by $z(t) := Q^\top x(t)$. Similarly, define $U = (u(t))_{t\in\mathbb{Z}}$ by $u(t) := Q^\top D\theta(t)$.

Since Q is orthogonal, X is bounded if and only if Z is bounded. Substituting $x(t) = Qz(t)$ into the coordinate form of (11) gives

$$Qz(t) - M'(\beta)Qz(t-1) = D\theta(t).$$

Left-multiplying by $Q^\top$, using $Q^\top M'(\beta)Q = \Lambda$, yields

$$z(t) - \Lambda z(t-1) = u(t).$$

Because Λ is diagonal, this separates into n scalar recurrences:

$$z_k(t) - \lambda_k z_k(t-1) = u_k(t), \qquad k = 1, \ldots, n.$$

By Theorem B.3, the kth scalar problem is well-posed on bounded histories if and only if $|\lambda_k| \neq 1$. Therefore the full system is well-posed if and only if no eigenvalue lies on the unit circle. Since $M'(\beta)$ is symmetric, its eigenvalues are real, so this is equivalent to excluding ± 1. □

Corollary B.5 (Transfer of the threshold $\beta_{\text{thresh}}(c)$)**.** *Under the hypotheses of Theorem 5.3, the threshold $\beta_{\text{thresh}}(c)$ from (6) remains the boundary between causal forward stability and noncausal bounded-history behavior:*

1. *If $\beta < \beta_{\text{thresh}}(c)$, then $\rho(M'(\beta)) < 1$. Hence for every bounded signal history Θ, there is a unique bounded solution, and each modal component depends only on present and past signal values.*

2. *If $\beta > \beta_{\text{thresh}}(c)$ and $\pm 1 \notin \sigma(M'(\beta))$, then bounded histories may still exist by Theorem B.4, but any component corresponding to an eigenvalue with modulus greater than 1 depends on future signal values. In this regime, the history equation may be well-posed as a two-sided object, but it no longer has a causal forward interpretation.*

Proof. If $\beta < \beta_{\text{thresh}}(c)$, then Theorem 5.3 gives $\rho(M'(\beta)) < 1$. Thus every eigenvalue satisfies $|\lambda_k| < 1$, and Theorem B.3 gives the causal formula

$$z_k(t) = \sum_{m=0}^{\infty} \lambda_k^m u_k(t-m).$$

So each modal component depends only on present and past inputs.

If $\beta > \beta_{\text{thresh}}(c)$, forward stability fails by Theorem 5.3. If, nevertheless, $\pm 1 \notin \sigma(M'(\beta))$, then Theorem B.4 still gives a unique bounded two-sided solution. For any mode with $|\lambda_k| > 1$, Theorem B.3 gives

$$z_k(t) = -\sum_{m=1}^{\infty} \lambda_k^{-m} u_k(t+m),$$

which depends on future inputs. Hence the two-sided history exists, but the forward causal interpretation is lost. □