Distributed Consensus Protocols in Adaptive Multi-Agent Systems

Aman Sinha, Class of 2013

May 2, 2013

Submitted to the
Department of Mechanical and Aerospace Engineering,
Princeton University
in partial fulfillment of the requirements
of Undergraduate Independent Work

Final Report

Professor Naomi E. Leonard
Professor Philip J. Holmes
MAE 442
Number of pages: 80
Color Images: Yes
Copy: Reader

Acknowledgements

I would like to thank numerous individuals who helped me put forth my best efforts in this study. First of all, I recognize Professor Leonard and George Young for introducing me to the world of networked systems and for guiding me throughout my first foray into the subject. Their mentorship, constructive suggestions, and patience helped stabilize the foundation of my research efforts and polish the end results. I would also like to thank Professor Holmes for taking the time to serve as a reader for my thesis and evaluate my work. Finally, I thank my colleagues and friends, including the 323 Society, for their feedback and camaraderie.

This thesis is dedicated to my family, for their unconditional love, support, and guidance.

Abstract

The analysis of noisy consensus dynamics in networks is of great interest to both advance the fundamental understanding of multi-agent systems in nature as well as create robust decentralized engineering systems. We develop a protocol that heuristically attempts to optimize metrics of consensus dynamics without explicitly measuring a network's global properties. Adopting the approach of utility maximization by nodes in a network, we allow nodes to modify connections with their neighbors over time. This results in a *locally adaptive* network: the global graph structure updates through the collective action of local changes, and no node has any knowledge of this evolution beyond its effects on the node's local environment. Our research focuses specifically on developing the form of this utility function to (heuristically) optimize network performance with respect to noisy consensus dynamics.

Beginning with a utility function inspired by economic and sociological models for network behavior, our analysis discovers the importance of coupling state and network dynamics to enhance consensus performance. Consequently, we develop the "perceived intelligence" coupling factor which creates a positive feedback between the state dynamics and network structure: nodes gravitate towards smart individuals who appear to be close to the final consensus state. Results indicate that this feedback reduces overshoot in the state dynamics and improves the convergence speed and robustness of consensus, but it induces heavy oscillations in network structure as individuals swing between smart individuals. Therefore, we sophisticate the model by introducing "intelligence history," a recursive estimation scheme for perceived intelligence that dampens the positive feedback, thereby reducing swings in network structure. With the addition of perceived intelligence and intelligence history, our protocol greatly outperforms the original utility model, especially when network costs are taken into account in the metrics of consensus performance. Overall, the protocol appears to be a very capable heuristic for maximizing consensus performance in the presence of noise, and it is easily adaptable to a variety of applications.

Contents

List of Figures

List of Symbols

$\bar{L}$ Reduced normalized Laplacian matrix

$\boldsymbol{\xi}(t)$ Noise for state dynamics

γ Variable that determines the length of history for perceived intelligence estimation

$\hat{f}_{ik}$ Estimation of perceived intelligence factor f_{ik} using history

$\hat{L}$ Normalized Laplacian matrix

λ_{min} Convergence speed for system

λ^*_{min} Convergence speed for a homogeneous complete BCC

λ^e_{min} Effective speed of convergence

λ^{cn}_{min} Cost-normalized effective speed of convergence

$\mathbf{h}$ Generic hidden mode vector

$\mathbf{h}^{c3}$ Symmetric zero-sum 3-cycle, the fundamental hidden mode

$\mathbf{x}$ Vector of states for all agents

$\mathcal{H}$ Space of hidden modes

$\mathcal{N}(g, h)$ Normal distribution with mean g and variance h

$\mathcal{N}_k$ Neighbors of node k

μ Parameter in cost function

$\deg(i)$ Unweighted undirected degree of node i (counts the neighbors in $\mathcal{N}_i$)

ρ Parameter in benefit function

σ_{ij} $a_{ij} + a_{ji}$

τ Parameter in benefit function

A Adjacency matrix

a_{ij} Edge directed from node i to node j

$B(\sigma_{ij})$ Benefit function for collaboration between nodes i and j

b_{ij} Dimensionless edge: $\frac{a_{ij}}{\rho}$

$C(d_i^{out})$ Cost function for node i

d_i^{in} Weighted indegree of node i

d_i^{out}	Weighted outdegree of node i
$d_0(\sigma)$	Nullcline
e_i^{out}	$\frac{d_i^{out}}{\rho}$
$e_0(\sigma)$	Dimensionless nullcline
f_{ik}	Node i's perception of node k's intelligence
H_2	H_2-norm for a system
H_2^*	H_2-norm for a homogeneous complete BCC
H_2^e	Effective H_2-norm of a system
H_2^{cn}	Cost-normalized effective H_2-norm
L	Laplacian matrix
m	$\mu\rho^2$
N	Number of agents in network
N_s	Number of nodes at or below which a homogeneous BCC is at least marginally stable
N_u	Number of nodes at or above which a homogeneous BCC is unstable
N_U^+	Number of nodes at or above which a symmetric BCC star is unstable
$NC(t)$	Network cost factor
OS	Overshoot for a state
P_i	Utility (i.e. payoff) function for node i
r	$\frac{\tau}{\rho^2}$
s_{ij}	$\frac{\sigma_{ij}}{\rho}$
t_s	Settling time for state dynamics
v_k	Variance of node k's state with its neighbors' states

1 Introduction

Characterizing the dynamics of multi-agent systems (MAS) and their performance towards a goal presents challenging problems for theoretical and computational analysis. Such systems consist of multiple interacting agents within an environment. Individual agents generally have a limited knowledge of their external environment and can interact in a limited fashion with other agents in the system. MAS are ubiquitous in nature and they dominate the realm of biological phenomena, from the macroscopic (flock migrations, honeybee swarms, and human social networks) to the microscopic (bacterial growth and intracellular phenomena). In addition, MAS have numerous scientific and engineering applications, particularly in domains where a centralized approach might be costly or intractable (automated disaster response or high-frequency trading in financial markets).

Within the MAS field, consensus dynamics studies networked agents working towards a common goal (i.e. tracking a reference signal or evading a predator). This particular focus has direct applications such as the optimization of distributed autonomous sensing networks or large-scale energy systems. The ability to reach consensus and the "quality" of this state (i.e. its stability, speed of convergence, and/or error with respect to a reference) depend on the network topology, the initial conditions of the system, the disturbances or noise inherent in the system, and the consensus protocol, which characterizes the network's evolution from initial conditions. Recent research efforts in consensus dynamics have studied various aspects of these factors, such as the effects of network topology on robustness of consensus to noise [16, 17] and the convergence speed of distributed consensus protocols [8]. Interesting results in the creation of robust network topologies have also been determined from the study of biological systems such as starling flocks [15].

Many studies in consensus dynamics focus on static network topologies due to greater analytical tractability. However, it is evident from even a cursory knowledge of biological systems, such as human social networks or flock migrations, that these types of network structures evolve or adapt with time. This evolution plays a large role in system performance with respect to consensus dynamics. Following the model of [15] and seeking inspiration from biological systems, current research focuses on the role of locally evolving topologies as a component of the individual consensus protocol for each agent (i.e. the distributed consensus protocol). In other words, we allow each agent to change the way it interacts with its neighbors regarding information retrieval and submission, thereby allowing the overall network structure to change with time. This makes the network *locally* adaptive. Locality of network evolution is of utmost importance for our research; this constraint preserves the sense of decentralization that differentiates MAS from more centralized schemes. From another perspective, imposing locality can also be considered realistic for many biological phenomena and engineering applications; it is likely that agents do not (or only intermittently) have the ability to measure global properties of the network, but they can determine local performance by interacting with neighbors. This locality limitation is not as crippling as it may initially seem. Even in physical systems, a neighbor need not be defined as an agent that is in close physical proximity to another agent. Abstracting the system as a graph of nodes and edges, neighbors are defined as nodes that are directly connected by an edge. Performance is driven by the network structure in topological space rather than application-specific metric spaces.

1.1 Objectives and Approach

In a framework where the distributed consensus protocol specifies only the behavior of individual agents with regard to their local environments, it is important to understand how local behavior affects global behavior. This study researches distributed consensus protocols that prescribe local network changes and analyzes the resulting effects on global performance with respect to noisy consensus dynamics. In other words, we examine noisy consensus dynamics in locally adaptive[1] networks using distributed consensus protocols. The analysis of this local/global correspondence is twofold: the first aspect involves defining the distributed consensus protocol and determining its effect on the (possibly steady-state) structure of the network; the second aspect involves analyzing the resulting structure with respect to metrics of global performance (i.e. robustness, stability, convergence speed, etc.).

The development of distributed consensus protocols has received considerable attention by various literatures including those of economics, control theory, game theory, and operations research. Current analysis combines the conventions and methodologies of control theory with models that resemble standard economic frameworks of utility maximization by individuals. Other approaches include system identification through machine learning [3] or the classic multi-armed bandit framework for researching explore/exploit situations [13]. The study of such approaches in the context of consensus dynamics is worth further research, but we do not focus on them here.

The remainder of this report is organized as follows. First, relevant background is presented regarding the motivation for the problem under consideration, followed by notation and preliminary material for analyzing noisy consensus dynamics. Section 2 presents a model for utility maximization, and Section 3 analyzes its performance with respect to consensus

[1] The term "locally adaptive" refers to the fact that the network structure evolves locally with time.

dynamics. Subsequently, we present a series of improvements to this utility model in Section 4 and characterize the performance of the updated models with respect to the original. The study is concluded with a review of the results, their implications, and open questions for future investigation.

1.2 Background

Comprehensive analysis of noisy consensus dynamics requires relevant background to introduce the problem and the associated measures for network performance. Section 1.2.1 motivates the overall approach presented in Section 1.1, and the subsequent material illustrates how to characterize the performance of multi-agent systems with respect to noisy consensus dynamics.

1.2.1 Motivation: Imperfect Information in Economic Markets

It is reasonable to ask whether or not the objectives presented in Section 1.1 are worth merit, i.e. whether the problem of reaching consensus through a distributed consensus protocol in the presence of noise is nontrivial. The most immediate answer (which is resoundingly affirmative) derives from economic literature. The model of an economic market is not very different from that considered here for consensus dynamics in MAS, particularly in the context of consensus protocols that involve local utility maximization. Indeed, an economic market consists of individuals maximizing utility (often subject to a budget constraint for consumers or a production constraint for producers). The First Welfare Theorem of Economics states that perfectly competitive markets[2] will reach Pareto-efficient[3] states in equilibrium (see [9] for a semi-rigorous proof). This result is often cited as an analytical confirmation for economist Adam Smith's famous "invisible hand" hypothesis on the self-regulatory nature of economic markets. However, the ideal conditions required by the First Welfare Theorem do not always hold in real economic markets. Namely, markets fail to reach Pareto-efficient states in the presence of imperfect competition (when individuals have market power to set prices such as in the case of a monopoly), externalities (interactions between individuals that are not adequately reflected by market prices), public goods (goods that are non-excludable and non-rivalrous such as fresh air or national defense), and imperfect information. The final factor, imperfect information, is particularly relevant to the noisy consensus dynamics problem presented in Section 1.1. As we show in Section 1.2.2, under certain conditions

[2]Perfect competition prescribes the inability of any individual or group to have the power to set prices or unfairly take advantage of others. In other words, there are no monopolies, oligopolies, or any types of information asymmetry between economic players.

[3]Pareto efficiency is the inability for any individual to increase utility without decreasing the utility of another. It is a local optimum in the space of resource allocations that maximize individuals' utility functions.

a network can be guaranteed to asymptotically reach consensus in the absence of noise, a reflection of the First Welfare Theorem. On the other hand, noise corrupts the information passed between neighbors and prevents such convergence. Whereas economies often resort to centralized measures such as government regulation to deal with factors that prevent Pareto efficiency, current research aims to retain the distributed, decentralized nature of the system. We study the ability of the system, as governed by the decentralized consensus protocol and network topology, to cope with noise corruption. From an economic perspective, this is equivalent to studying a market's innate ability to deal with imperfect information without intervention by any other agent. In this way, the objectives presented in Section 1.1 set up a nontrivial and rich problem.

1.2.2 Model for Consensus Dynamics in Multi-Agent Systems

The model presented here is widely used throughout control-theoretic literature. Our nomenclature derives from that presented in [15, 16, 17]. We consider a system of N agents such that the state of the network is represented by $\mathbf{x} = [x_1, x_2, ..., x_N]^T \in \mathbb{R}^N$. Then the system is in consensus when $\mathbf{x} = \alpha \mathbf{1}_N$, where $\mathbf{1}_N = [1, 1, ..., 1]^T \in \mathbb{R}^N$ and $\alpha \in \mathbb{R}$. Alternatively, we can define the subspace orthogonal to $\mathbf{1}_N$ as $\Pi = I_N - \frac{1}{N}\mathbf{1}_N\mathbf{1}_N^T$, where I_N is the identity matrix. The consensus state satisfies $\Pi\mathbf{x} = 0$.

Graph Structure and Nomenclature The network structure is defined by a directed graph, $G = (V, E, A)$, where $V = \{1, 2, ..., N\}$ is the set of nodes that represent agents, $E \subseteq V \times V$ is the set of edges that represent neighborly connections, and $A \in \mathbb{R}^{N \times N}$ is the weighted adjacency matrix defining the weight of each neighborly connection. Specifically, every entry a_{ij} in A is nonnegative, and $a_{ij} > 0 \iff (i, j) \in E$. Directed edges point from receivers of information to senders: a_{ij} points from node i (the receiver) to node j (the sender). In other words, the directed edges leaving node i define agent i's neighbors (agents that agent i observes) and a_{ij} is the weight with which agent i gives to neighbor j's information. We prescribe that there exist no self-cycles: $a_{ii} = 0 \ \forall \ i \in V$. The weighted outdegree of node i is defined as $d_i^{out} = \sum_j a_{ij}$ and the weighted indegree by $d_i^{in} = \sum_j a_{ji}$. Finally, a path in G represents a directed sequence between nodes along directed edges. The graph G is called connected if there exists some node that is reachable by a path from any other node. We prescribe that networks are connected for further analysis, as this has great ramifications on system dynamics.

The dynamics of network evolution are governed by the Laplacian of G, defined as $L = D - A$, where D is a diagonal matrix with d_i^{out} as the i^{th} diagonal entry. Note that the row sums of L are 0, so $\mathbf{1}_N$ is an eigenvector of L with eigenvalue 0 ($L\mathbf{1}_N = 0$). Furthermore, 0 is a simple eigenvalue (multiplicity of 1) if and only if G is connected (Theorem 4 of [1]). Application of the fiendishly simple Gershgorin Circle Theorem (Theorem A.1 of Appendix A) shows that all eigenvalues λ of L satisfy $\mathrm{Re}\,\lambda \geq 0$. To normalize the magnitude of the eigenvalues with respect to a_{ij}, we define a normalized Laplacian ($\hat{L}$) as one in which the diagonal entries are either 0 or 1. Therefore, $\hat{L}_{ij} = L_{ij}/d_i^{out}$. The intuition behind this normalization is that only the relative magnitudes of a_{ij} should affect the dynamics in $\mathbf{x}$. This implies that an isolated network of highly social individuals could have the same $\mathbf{x}$ dynamics as one with more introverted individuals; a uniform scaling in a_{ij} has no effect on $\dot{\mathbf{x}}$. The majority of this study is focused on evaluating techniques that dynamically change the edge weights a_{ij} with time along with the state ($\dot{L} \neq 0$).[4]

Noisy Consensus Dynamics and Convergence Speed Noise corrupts the information flow between along edges either from output noise from the sender, measurement noise by the receiver, disturbances, or a combination of all corrupting phenomena. Following [16], we consider uniform white noise across all nodes. Therefore, we define the network evolution for a static network structure ($\dot{L} = \dot{\hat{L}} = 0$) as:

$$\dot{\mathbf{x}}(t) = -\hat{L}\mathbf{x}(t) + \boldsymbol{\xi}(t) \tag{1.1}$$

where $\boldsymbol{\xi}(t) \in \mathbb{R}^N$ is a random signal satisfying the usual properties for white noise: $E[\boldsymbol{\xi}(t)] = 0$, $E[\boldsymbol{\xi}(t)\boldsymbol{\xi}^T(\tau)] = \eta^2 I_N \delta(t-\tau)$, and $E[\mathbf{x}(0)\boldsymbol{\xi}^T(\tau)] = 0$, where the autocorrelation intensity of noise is prescribed by η.

The system defined above is only marginally stable: one eigenvalue is 0 and all others are below 0 for the matrix $-\hat{L}$.[5] Because the marginally stable state is the consensus state, it is conceivable that we may only need to look only at the system dynamics along the space spanned by Π, the dispersion of the system from equilibrium.[6] Namely, we define a (non-unique) matrix Q with rows forming an orthonormal basis for Π: $Q\mathbf{1}_N = 0$, $QQ^T = I_{N-1}$, and $Q^TQ = \Pi$. The dispersion is then defined as $\sqrt{\mathbf{y}^T(t)\mathbf{y}(t)}$, where $\mathbf{y} := Q\mathbf{x}$. This results

[4] Note that $\dot{L} \neq 0 \not\Rightarrow \dot{\hat{L}} \neq 0$. One simple example is the case of edge weights that scale uniformly over the entire graph ($\dot{a}_{ij} = c$, where c is a constant).

[5] As noted earlier, we only consider connected network topologies in this study, which is why 0 is a simple eigenvalue. As we will later show, unconnected network topologies correspond to non-collaborative systems in which consensus is virtually impossible.

[6] The dynamics projected onto $\mathbf{1}_N$ characterize the mean state of the system, which is marginally stable. Current research does not analyze the ability of the system to reach a certain value at the consensus state, but rather studies the performance of the system in purely reaching and/or maintaining a consensus state. This analysis can be generalized to the former scenario in numerous ways. One method would be to augment the network with "leader" nodes who track an exogenous signal.

in the reduced system dynamics:

$$\dot{\mathbf{y}}(t) = -Q\hat{L}Q^T\mathbf{y}(t) + Q\boldsymbol{\xi}(t) = -\bar{L}\mathbf{y}(t) + Q\boldsymbol{\xi}(t) \tag{1.2}$$

where $\bar{L}$ is the reduced normalized Laplacian of the system [16]. Since Q is not unique, neither is $\bar{L}$. However, for any second parametrization $\mathbf{y}'$ and Q', we find that $\mathbf{y}'^T\mathbf{y}' = (Q'\mathbf{x})^T Q'\mathbf{x} = (Q'Q^T\mathbf{y})^T Q'Q^T\mathbf{y} = \mathbf{y}^T\mathbf{y}$, so the dispersion is the same [16]. Importantly, the eigenvalues of $\bar{L}$ are the nonzero eigenvalues of $\hat{L}$ for a connected graph, whereby the dispersion dynamics are stable [16]. The speed of convergence is then dominated by the slowest eigenmode, so we define the convergence speed as the real part of the smallest eigenvalue in $\bar{L}$. As this asymptotic stability is a key result for further analysis, we review the proof here [16]:

Theorem 1.1. *The eigenvalues of $\bar{L}$ are the nonzero eigenvalues of $\hat{L}$ for a connected graph.*

Proof. Define the orthogonal change of basis:

$$V = \begin{bmatrix} Q \\ \frac{1}{\sqrt{N}}\mathbf{1}_N^T \end{bmatrix}$$

Then the system dynamics (i.e. the eigenvalues) of matrix $V\hat{L}V^T$ must be the same as those of $\hat{L}$. However, the new matrix has the following block triangular structure:

$$VLV^T = \begin{bmatrix} \bar{L} & \mathbf{0}_{N-1} \\ \frac{1}{\sqrt{N}}\mathbf{1}_N^T\hat{L}Q^T & 0 \end{bmatrix}$$

so its eigenvalues are the union of the eigenvalues of $\bar{L}$ and zero. Since zero is a simple eigenvalue of $\hat{L}$, $\bar{L}$ has the nonzero eigenvalues of $\hat{L}$. □

Robustness of Consensus Noise perturbs the system and prevents it from reaching consensus otherwise guaranteed by the reduced system's asymptotic stability. The robustness of consensus with respect to noise can be defined through different types of norms such as the H_2- and H_∞-norms. Whereas the H_2-norm characterizes the output energy for a system due to a unit level disturbance, the H_∞-norm characterizes the maximum magnitude of response (or maximum amplification) of noise possible in a system. Current research will focus on the H_2-norm (abbreviated as H_2).

Consider a generic asymptotically stable contious time linear-time-invariant multiple-input-multiple-output (CTLTI-MIMO) dynamical system with state $\mathbf{x}$, (disturbance) inputs $\mathbf{w}$, and output $\mathbf{y}$:[7]

$$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{w} \tag{1.3a}$$

$$\mathbf{y} = C\mathbf{x} \tag{1.3b}$$

$$G(s) = C(sI - A)^{-1}B \tag{1.3c}$$

$$\mathbf{g}(t) = \mathcal{L}^{-1}\left\{G(s)\right\} = Ce^{At}B \tag{1.3d}$$

where Equation 1.3c is the transfer function of the system, and Equation 1.3d is its corresponding inverse Laplace transform (the impulse response). Then, the square of the H_2-norm (H_2^2) can be defined in the following equivalent ways using the transfer function or the impulse response:

$$\begin{aligned} H_2^2 &= \frac{1}{2\pi}\operatorname{tr}\left[\int_{-\infty}^{\infty} G(j\omega)G^*(j\omega)d\omega\right] \\ &= \operatorname{tr}\left[\int_0^{\infty} \mathbf{g}(t)\mathbf{g}^T(t)dt\right] \\ &= \operatorname{tr}\left[B^T W_O B\right], W_O = \int_0^{\infty} e^{A^T t}C^T C e^{At}dt \\ &= \frac{1}{2\pi}\operatorname{tr}\left[\int_{-\infty}^{\infty} G^*(j\omega)G(j\omega)d\omega\right] \\ &= \operatorname{tr}\left[\int_0^{\infty} \mathbf{g}^T(t)\mathbf{g}(t)dt\right] \\ &= \operatorname{tr}\left[CW_C C^T\right], W_C = \int_0^{\infty} e^{At}BB^T e^{A^T t}dt \end{aligned} \tag{1.4}$$

where the conversion from time domain to frequency domain is permissible by Parseval's Relation. W_O and W_C are the observability and controllability Gramians, which are the solutions to the following dual Lyapunov equations:[8]

$$A^T W_O + W_O A + C^T C = 0 \tag{1.5a}$$

$$AW_C + W_C A^T + BB^T = 0 \tag{1.5b}$$

The reduced consensus dynamics system can be rewritten in the following way (for static

[7] The $\mathbf{x}$ and $\mathbf{y}$ variables here are named as such merely as dictated by convention. They do not correspond to the network state and consensus dispersion state introduced for the consensus dynamics problem.

[8] In discrete-time systems, these equations are slightly modified: $A^T W_O A - W_O + C^T C = 0$ and $AW_C A^T - W_C + BB^T = 0$.

network structures):

$$\dot{\mathbf{y}} = -\bar{L}\mathbf{y} + Q\boldsymbol{\xi} \tag{1.6a}$$

$$\mathbf{z} = I_{N-1}\mathbf{y} \tag{1.6b}$$

Then, the H_2-norm is $\sqrt{\operatorname{tr}\left[I_{N-1}PI_{N-1}\right]} = \sqrt{\operatorname{tr} P}$, where $\bar{L}P + P\bar{L}^T - I_{N-1} = 0$. However, it is also possible to arrive at this result in a slightly more direct fashion that gives meaning to P [16]. Loosely speaking, robustness is measured by the "energy" of the output, so it reasonable to consider the expected value of the output over all time due to a disturbance: $\lim_{t\to\infty} E[\|\mathbf{y}\|] = \sqrt{\operatorname{tr}\Sigma_{ss}}$, where $\Sigma_{ss} = \lim_{t\to\infty}\Sigma(t) := \lim_{t\to\infty} E[\mathbf{y}(t)\mathbf{y}^T(t)]$. Calculating $\dot{\Sigma}(t)$ and setting it to 0 as $t \to \infty$ yields the result that $\bar{L}\Sigma_{ss} + \Sigma_{ss}\bar{L}^T - 2\eta^2 I_{N-1} = 0$ (see Lemma A.3), whereby $\Sigma_{ss} = P$ when $\eta^2 = \frac{1}{2}$. It is important to note that these statements are valid only for LTI systems or static network structures. Therefore, they are most applicable to the current study under conditions where the system has reached a steady-state structure. However, we will also analyze the evolution of the "snapshot" H_2-norms of the network as its structure changes with time. Robust systems minimize the H_2-norm.

2 Framework for Network Evolution

This section outlines the overall framework upon which our current research on distributed consensus protocols is conducted. Namely, we present a basic model for network evolution through utility maximization. Equation 1.1 summarizes the *state* dynamics, but does not prescribe a model for network evolution ($\dot{L}$). When considering network evolution through utility maximization, it is obvious that the specific choice of utility function may drastically affect behavior. We turn to nature for inspiration regarding the form of utility in a collaborative network. In [5], the evolution of network structure is studied according to a specially designed utility maximization protocol that builds upon economic and sociological models of cooperative interaction in [6, 10]. We present this model in its original form and also derive a dimensionless form that yields insights into how the network dynamics scale relative to state dynamics. Then, we briefly survey key results from [5] for network equilibria.

2.1 Utility Model

The utility maximization problem is best posed in terms of resource allocation by individual agents. Each neighborly connection consists of a collaboration, and each agent i invests a_{ij} into its collaboration with agent j. As in canonical economic models that maximize profit (revenue subtracted by cost), each collaboration has a payoff which is given by a benefit function subtracted by a cost function. The functional form of benefit and cost are the main ideas to be extracted from sociological models, as shown below.

The benefit of a collaboration is determined to be a sigmoidal function of the variable $\sigma_{ij} = a_{ij} + a_{ji}$ while the cost to an individual is a superlinear function of the individual's weighted outdegree, d_i^{out}. This can be motivated using the following model situation: consider a group of students collaborating on a tough homework problem in a classroom. The students are only allowed to leave the room when the entire group has reached a consensus on the answer. Furthermore, students are only allowed to collaborate with other students in the room who are sitting in desks adjacent to their own. Then, it stands to reason that the benefit of an individual collaboration must saturate at some level of investment, whereas the cost to a student of paying attention to his neighbors will increasingly require greater concentration, energy, time, etc. Both students benefit from a collaboration, even if the collaboration is highly skewed (smart students still need lesser-abled students to reach the consensus with everyone else). Hence, the benefit is symmetric in the pair (i, j). This model captures realistic phenomena such as the "inefficiency of small investments, saturation of benefits at high investments, as well as additional costs incurred by overexertion of personal

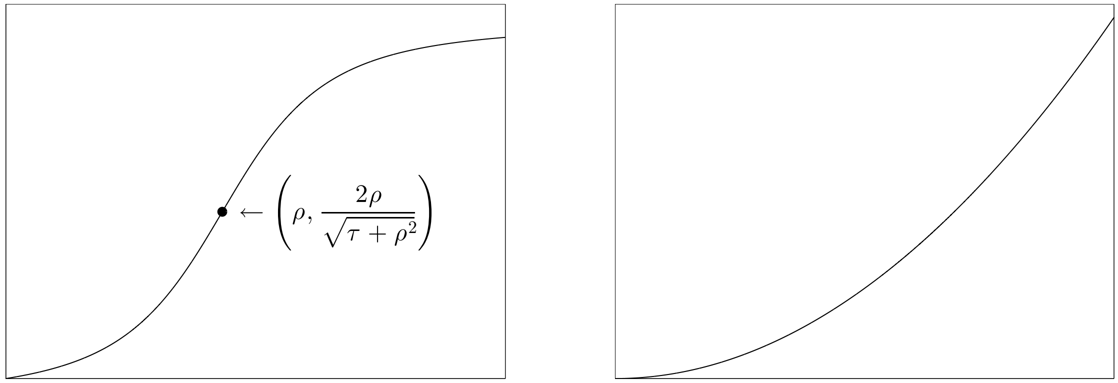

(a) $B(\sigma_{ij})$ vs. σ_{ij} (b) $C(d_i^{out})$ vs. d_i^{out}

Figure 2.1. (a) Benefit and (b) Cost functions for the utility model.

resources" [5]. A summary of the utility for each agent is:

$$P_i = \sum_{j \neq i} B(\sigma_{ij}) - C\left(d_i^{out}\right) \tag{2.1}$$

where B is sigmoidal and C is superlinear. The specific functional forms are further constrained by the fact that $B(0) = C(0) = 0$. Also, we intuit that the x-intercept should lie below the point of inflection of the sigmoid's S-curve to allow for the point of inflection (the point above which the marginal benefit starts decreasing) to play a part in the dynamics of collaboration. Example functional forms are given in [5]:

$$B(\sigma_{ij}) = \frac{2\rho}{\sqrt{\tau+\rho^2}} + \frac{2(\sigma_{ij}-\rho)}{\sqrt{\tau+(\sigma_{ij}-\rho)^2}} \tag{2.2a}$$

$$C(d_i^{out}) = \mu(d_i^{out})^2 \tag{2.2b}$$

for parameters ρ, τ, and μ. The point of inflection occurs at $\sigma_{ij} = \rho$, so larger ρ shifts the point of inflection to the right. Therefore, we require $\rho > 0$ for the point of inflection to play a part in the dynamics. Larger τ flattens the sigmoid: it approximates a line as $\tau \to \infty$. The parameter μ governs the cost-benefit ratio for a given network. Figure 2.1 shows representative benefit and cost functions for this utility model.

Each agent maximizes utility at any given time by performing gradient ascent, resulting in the following system for (non-reduced) consensus dynamics:[9]

$$\dot{\mathbf{x}}(t) = -\hat{L}(t)\mathbf{x}(t) + \boldsymbol{\xi}(t) \tag{2.3a}$$

[9] The state dynamics are linear with static $\hat{L}$. Now both the state dynamics and network dynamics are nonlinear.

$$\dot{a}_{ij}(t) = \frac{\partial}{\partial a_{ij}} P_i(t) \quad \forall (i,j) \in E \tag{2.3b}$$

Importantly, Equations 2.1 and 2.3b use non-normalized edge-weights. Otherwise, we would have $d_i^{out} \in \{0,1\}\ \forall\ i \in V$ and therefore only one nonzero value for $C(d_i^{out})$. This would result in unrealistic dynamics for the utility maximization model prescribed. Thus, whereas $\dot{\mathbf{x}}$ depends only on $\hat{L}$ to isolate the effects of the absolute magnitudes in a_{ij}, Equation 2.3b is a model for $\dot{L}$ and, as a result, the absolute magnitudes of a_{ij} affect the dynamics of network structure evolution. As a final note, the number of agents N has effects on both the state dynamics (through the eigenvalues of $\hat{L}$) and network structure dynamics (to be shown in Section 3).

2.1.1 Dimensionless Form and Network Timescale

We now use dimensional analysis to relate the dynamics of Equation 2.3b to the state dynamics. The state dynamics evolve in time proportional to the inverse of the smallest eigenvalue λ of $\hat{L}$, which is guaranteed to satisfy $\operatorname{Re}\lambda \in [0,2]$ (see Corollary A.2). If we consider the functional form for $B(\sigma_{ij})$, then we see that the units of σ_{ij} are the same as those of ρ and $\sqrt{\tau}$, rendering $B(\sigma_{ij})$, $C(d_i^{out})$, and P_i dimensionless. Then, assigning units of Col (for collaboration) to a_{ij}, Equation 2.3b requires that time be proportional to Col^2. Following the form of the classic Buckingham Π Theorem [4], we can find a convenient basis for the kernel of the matrix of dimensions (the Π-products or dimenionless parameters), provided that we properly choose the form with which to scale a_{ij}.[10] Because the network dynamics must occur about the inflection point of the sigmoidal benefit curve, the relevant scale in this case is ρ. In other words, ρ essentially shifts the location of the sigmoid in σ_{ij}, thereby changing the characteristic value for the sum $a_{ij} + a_{ji}$.

Using this scaling, we define the following dimensionless variables:

$$b_{ij} = \frac{a_{ij}}{\rho},\ s_{ij} = \frac{\sigma_{ij}}{\rho},\ e_i^{out} = \frac{d_i^{out}}{\rho}$$

and derive the following dimensionless components of utility:

$$B(s_{ij}) = \frac{2}{\sqrt{r+1}} + \frac{2(s_{ij}-1)}{\sqrt{r+(s_{ij}-1)^2}} \tag{2.4a}$$

[10]This caveat is similar to considering the manner in which to scale pressure in the dimensionless form of the Navier-Stokes equations. Depending on the flow regime, pressure may scale with viscous forces or inertial forces. Choosing the wrong form can result in terms blowing up or vanishing in the dimensionless formulation.

$$C(e_i^{out}) = m(e_i^{out})^2 \tag{2.4b}$$

$$r = \frac{\tau}{\rho^2},\ m = \mu\rho^2 \tag{2.4c}$$

Finally, nondimensionalizing Equation 2.3b yields ρ^2 as the relevant timescale for characterizing network structure dynamics. Thus, whereas the state dynamics are limited to a timescale $\geq \frac{1}{2}$, the network dynamics can be conceivably faster based on the value of ρ. Of course, the network dynamics near equilibrium will not necessarily occur at this timescale. They are governed by the distance away from ρ at which equilibrium points lie, and, as we show in Section 3.3.1, equilibrium perturbations are governed by the timescale $\mu^{-1} = \rho^2/m$.

Importantly, Equation 2.2 is equivalent to Equation 2.4, but, as shown here, the latter provides a clearer relationship between state and network dynamics, which encompasses a large portion of this study.

2.2 Review of Previous Results for Network Dynamics

This section reviews results regarding the dynamics of network evolution, i.e. Equations 2.2 and 2.3b, which was conducted in [5]. To reiterate, [5] does not study evolution of a state vector $\mathbf{x}$ along with network structure; it only studies the evolution of the network structure itself. However, this analysis is important as it helps us derive key properties regarding the steady-state structure of networks with this utility model, which has immediate ramifications on the state dynamics.

The steady-state behavior can be characterized by the coordination of investments a_{ij}. We define the pair (i, j) as a *bidirectional* connection when $a_{ij} \neq 0 \iff a_{ji} \neq 0$ and define a *bidirectionally connected component* (BCC) as a connected sub-graph of the network composed solely of bidirectional connections. Similarly, a one-sided link is deemed *unidirectional*. In general, steady states will consist of both unidirectional and bidirectional connections, although the latter type is more common (shown below). In addition, links that begin with nonzero values but vanish to zero at steady-state are *vanished* links.[11] Qualitative behavior is summarized in Figure 2.2, the result of a simulation carried out using an adaptive forward-Euler method with $a_{ij}(t_0) = \mathcal{N}(1, 10^{-28})$, $\rho = 0.65$, $\tau = 0.1$, and $\mu = 1.5$ [5].

The following two properties regarding the coordination of investments within a BCC yield important ramifications regarding steady-state network structure, so we review the proofs from [5]:

[11] Numerically, a link a_{ij} is said to have vanished when its value goes below a certain threshold and $\dot{a}_{ij} < 0$. In this study, we use a threshold of 10^{-10}. Note that a vanished link may "unvanish" if $\dot{a}_{ij} > 0$.

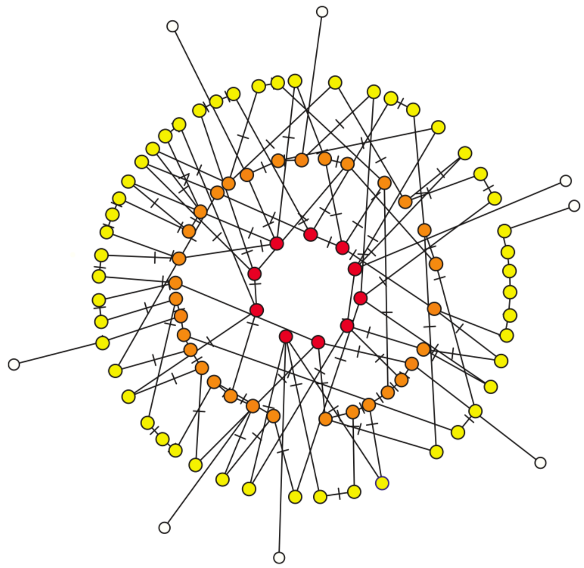

Figure 2.2. Representative simulation of steady-state network structure [5]. Lines between nodes indicate connections, with a dash to indicate the relative magnitude of a_{ij} to a_{ji}; a dash closer to node i indicates $a_{ji} > a_{ij}$. Individuals with higher utility are placed in a more central location and are shaded darker. These individuals are characterized by a relatively high number of collaborations in which they provide minimal effort, thereby receiving high benefits for little cost. Fringe individuals (in white) collaborate with others through unidirectional links.

Proposition 2.1. *Within a BCC at equilibrium, $d_i^{out} = c$ where $c \in \mathbb{R}_+$, or all agents make the same total investment.*

Proof. By symmetry,

$$\frac{\partial B(\sigma_{ij})}{\partial a_{ij}} = \frac{\partial B(\sigma_{ij})}{\partial a_{ji}} := B'(\sigma_{ij})$$

$$\frac{\partial C(d_i^{out})}{\partial a_{ij}} = \frac{\partial C(d_i^{out})}{\partial a_{ik}} := C'(d_i^{out})$$

Stationarity of a bidirectional link requires $\dot{a}_{ij} = \dot{a}_{ji} = 0$ as well as $a_{ij} \neq 0$, which implies $B'(\sigma_{ij}) = C'(d_i^{out}) = C'(d_j^{out})$. Then, because C' is injective, $d_i^{out} = d_j^{out}$ and the pair (i, j) exhibit the same total investment. Because a BCC is composed solely of bidirectional links, this argument iterates across all pairs of nodes within a BCC. $\square$

Proposition 2.2. *At equilibrium, a bidirectional link between the pair (i, j) within a BCC receives one of up to two possible values for investment ($\sigma_{ij} \geq \rho$ and $\sigma_{ij} \leq \rho$). If a BCC is locally stable to arbitrary perturbations in a_{ij}, each node can have at most one bidirectional link with σ_{ij} at the lower of the two possible values.*

Proof. By Proposition 2.1, $B'(\sigma_{ij})$ is the same for all pairs (i, j) within a steady-state BCC.

If this were not true, a node i could rearrange its investments to increase utility, violating stationarity. Due to the sigmoidal form of $B(\sigma)$, there are at most two locations for this value of $B'(\sigma_{ij})$, one above and one below the inflection point (or they may both lie on the inflection point).

It is shown in [5] that local stability of a BCC to perturbations in a_{ij} has the following necessary conditions for a single bidirectional link:

$$C''(d_i^{out}) > 0 \text{ and } 2B''(\sigma_{ij}) - C''(d_i^{out}) < 0 \tag{2.6}$$

and the following necessary conditions for every pair of bidirectional links (i, j) and (i, k) in the BCC:

$$B''(\sigma_{ij})B''(\sigma_{ik}) > C''(d_i^{out})\left(B''(\sigma_{ij}) + B''(\sigma_{ik})\right) \tag{2.7}$$

The quadratic form for C makes $C'' > 0$ always. Enumerating through the possible combinations of $B''(\sigma_{i*}) > 0$ and $B''(\sigma_{i*}) < 0$ yields the result that node i can have at most one link with σ_{i*} below the inflection point. □

There are a few interesting remarks about the above propositions. First, Proposition 2.2 shows the importance of the inflection point ($\sigma_{ij} = \rho$) of the sigmoidal benefit curve in the stability of network dynamics. Importantly, $\rho = \arg\max B'(\sigma_{ij})$, and $B'(\sigma_{ij})$ is monotonically decreasing in $|\sigma_{ij} - \rho|$. Thus, for a fixed level investment by node i, a withdrawal in a_{ji} makes node j more (less) attractive when σ_{ij} is above (below) ρ. This behavior accurately models modes of social interaction, as observed in [2] and (humorously) described in [11, 12].

Additionally, we emphasize that the local stability conditions discussed in Proposition 2.2 are necessary but not sufficient conditions for stability in arbitrary perturbations of a_{ij}. We show in Section 3.2 that some BCC's are at best marginally stable, even when Equations 2.6 and 2.7 hold true. This results from hidden modes of collaboration that are not captured by network dynamics.

Finally, it would appear that the steady-state structure for a BCC is severely limited since d_i^{out} and σ_{ij} are constrained to one or two values respectively. However, this degree of homogeneity does not preclude heterogeneity in a_{ij} within a BCC. To get an intuitive notion for heterogeneity in BCC links (as shown in Figure 2.2), note that Propositions 2.1 and 2.2 constrain the *total* investment for each node but only constrain the benefit for an *individual* collaboration. Thus, the benefit per node scales with the (unweighted) nodal degree, yielding higher benefits and overall utility for nodes with higher degree centrality (indegree and outdegree centrality are equivalent in a BCC). Additionally, degree-central nodes must nec-

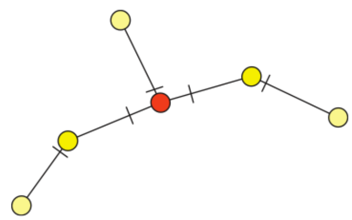

Figure 2.3. Heterogeneity in BCC links [5]. This heterogeneity arises from variations in degree centrality. Degree central nodes pay little attention to their neighbors, but receive lots of attention from these neighbors. Because benefits scale with degree centrality while nodal costs remain fixed for all nodes, individual payoffs scale with degree centrality.

essarily place less investment into an individual collaboration at a stationary state than less degree-central nodes due to the conservation of total investment for each node. Therefore, investments "flow" towards areas of higher average degree centrality, as shown in Figure 2.3.

This observation raises an important question regarding link heterogeneity in BCC's which are homogeneous in (unweighted) degree, such as a cycle BCC (nodal degree of 2) or a complete BCC (nodal degree of $N-1$). We show in Section 3.2 that heterogeneity in a_{ij} is possible in degree-homogeneous BCC's for $N > 2$ due to the hidden modes of collaboration mentioned above, and this characteristic plays an important part in the state dynamics of BCC's. First, we must present a more natural basis with which to represent the network dynamics.

3 Analysis of Network and State Dynamics

In this section, we analyze network and state dynamics using the model presented in Section 2 for $\dot{L}$. Having illustrated results from previous studies to motivate our current research, we present further analysis regarding network evolution and its ramifications on state dynamics. This leads to important conclusions regarding the shortcomings of the current utility model with respect to consensus dynamics.

3.1 Network Dynamics in (σ, d) Space

Although the model from [5] poses network dynamics in terms of a_{ij} (Equation 2.3b), it is evident from the functional forms of $B'(\sigma_{ij})$ and $C'(d_i^{out})$ that network evolution only depends on the variables σ_{ij} and d_i^{out}, a basis that we will henceforth abbreviate as the (σ, d) space. Note that a collaboration between nodes i and j has two sets of (σ, d) dynamics: (σ_{ij}, d_i^{out}) and (σ_{ij}, d_j^{out}), but these dynamics are just trajectories in (σ, d) space and can thus be plotted together. Therefore, the most useful aspect of analyzing the problem in this space is the fact that we can conveniently illustrate the network evolution as a collection of trajectories in two dimensions; the exact number of trajectories scales with network size and depends on the graph structure.

The nullclines for network dynamics are defined as the manifolds along which $\dot{a}_{ij} = 0$. We define the pair (σ_0, d_0) as a point along this manifold for some arbitrary pair of nodes (i, j). Solving $\dot{a}_{ij} = 0$ yields the nullcline as an algebraic function $d_0(\sigma)$ in (σ, d) space (or $e_0(s)$ in dimensionless form):[12]

$$d_0(\sigma) = \frac{\tau/\mu}{\left[\tau + (\sigma - \rho)^2\right]^{\frac{3}{2}}} \ , \quad e_0(s) = \frac{r/m}{\left[r + (s-1)^2\right]^{\frac{3}{2}}} \tag{3.1}$$

The relevant dimensionless form can be found by just replacing the variables $(d, \sigma, \tau, \mu, \rho)$ with $(e, s, r, m, 1)$ according to Section 2.1.1.[13] This nullcline is shown in Figure 3.1 and is marked by two special points separating different regimes for network dynamics near equilibrium.

The first special point is $\sigma = \rho$, which is a characteristic value when considering local stability against perturbations in σ or d from equilibrium. As noted in Figure 3.1, $\dot{a}_{ij} > 0$ underneath the nullcline and $\dot{a}_{ij} < 0$ above the nullcline. Also, we know that $\partial d_i / \partial a_{ij} = \partial \sigma_{ij} / \partial a_{ij} = 1$.

[12] For clarity, we will omit writing $*_{ij}$ for expressions that are for arbitrary pairs of nodes.

[13] Due to this similarity, we will often present results in dimensional form for consistency with previous results.

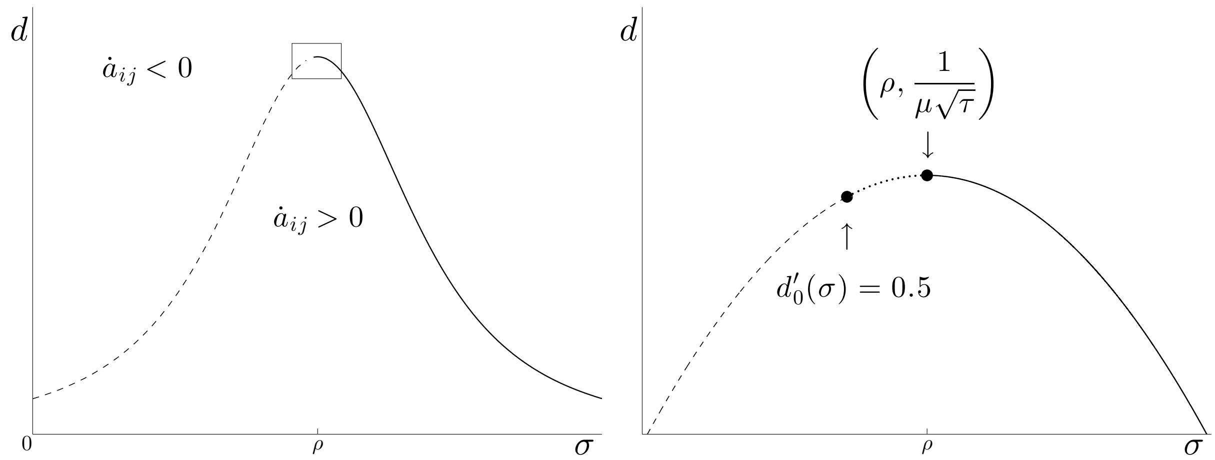

(a) Nullcline

(b) Zoomed-in view of nullcline near $\sigma \approx \rho$

Figure 3.1. The nullcline in (σ, d) space: $d_0(\sigma)$. For $\sigma \geq \rho$, the dynamics are locally stable to perturbations in σ and d. For $d_0'(\sigma) < 0.5$ (trivially satisfied for $\sigma \geq \rho$), a single bidirectional connection between nodes i and j is locally stable to perturbations in a_{ij}. ($d_0'(\sigma) := \frac{d}{d\sigma} d_0(\sigma) = \frac{d}{ds} e_0(s) := e_0'(s)$).

Thus, local perturbations in d always tend to move the system back towards the nullcline, although not necessarily back to the same point on the nullcline due to the coupling of the numerous (σ, d) pairs. In contrast, local perturbations in σ tend to move the system back towards the nullcline only when $\sigma \geq \rho$. Thus, $d_0(\sigma)$ exhibits saddle-like behavior for $\sigma < \rho$ under local perturbations in σ and d.

The second special point characterizes regions of stability to local perturbations in a_{ij} from equilibrium for a single bidirectional collaboration. By Proposition 2.2, a bidirectional collaboration is stable against local perturbations in a_{ij} when $2B''(\sigma) < C''(d)$ (Equation 2.6). This corresponds to $0.5 > d_0'(\sigma)$. We can now complement the proof of this result in [5] with a graphical argument. The magnitude and direction of $\dot{a}_{ij} = \partial P_i / \partial a_{ij} = B'(\sigma_{ij}) - C'(d_i^{out})$ depends on the vertical distance of the state from the nullcline: $d_i^{out} - d_0(\sigma_{ij})$. Namely, $\dot{a}_{ij} = -2\mu \left[d_i^{out} - d_0(\sigma_{ij})\right]$. Consider a point on the nullcline $(\sigma_0, d_0(\sigma_0))$. Then a small perturbation (δ_1, δ_2) from this point to $(\sigma_0 + \delta_1, d_0(\sigma_0) + \delta_2)$ results in the following linearized dynamics:

$$(\dot{\Delta a_{ij}}) = -2\mu \left[(d_0(\sigma_0) + \delta_2) - d_0(\sigma_0 + \delta_1)\right] \approx -2\mu \left[\delta_2 - \delta_1 d_0'(\sigma_0)\right] \tag{3.2}$$

The perturbation δ in a_{ij} from equilibrium corresponds to the perturbation (δ, δ) for agent i and $(\delta, 0)$ for agent j in (σ, d) space. We already know that the collaboration is stable when $\sigma \geq \rho$ (i.e. $d_0'(\sigma) \leq 0$), and this is also evident from Equation 3.2. For $d_0'(\sigma) > 1$, both perturbed points lie below the nullcline, so the collaboration is unstable. For $0 < d_0'(\sigma) \leq 1$,

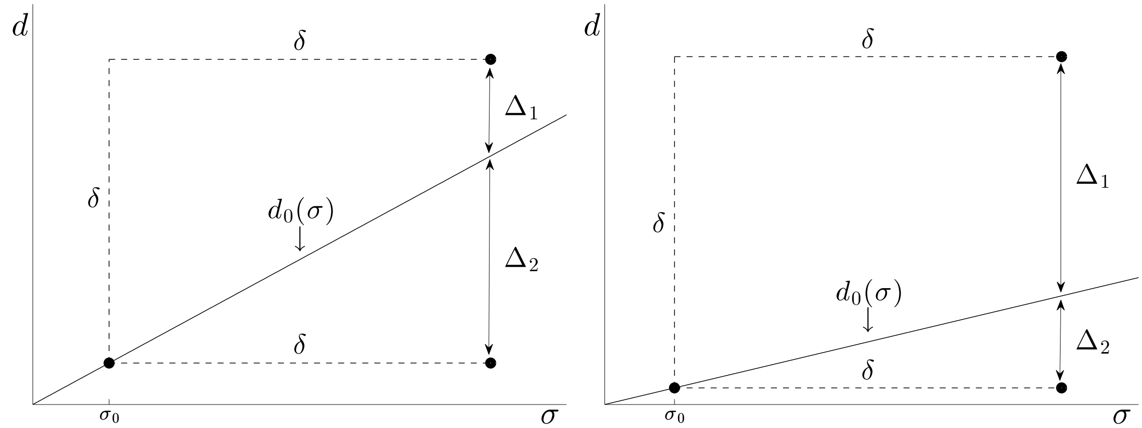

(a) Unstable since $\Delta_1 < \Delta_2$ (b) Stable since $\Delta_1 > \Delta_2$

Figure 3.2. Local stability analysis of a single bidirectional collaboration for $0 < d_0'(\sigma) \leq 1$. These pictures show a zoomed-in view of small perturbations from equilibrium, wherein $d_0(\sigma)$ appears approximately linear.

the perturbed dynamics are dominated by the perturbation that moved farthest away from the nullcline (for small δ). The relative magnitudes of $(\Delta\dot{a}_{ij})$ and $(\Delta\dot{a}_{ji})$ can be determined from a first-order approximation of Equation 3.2:

$$\left|\frac{(\Delta\dot{a}_{ij})}{(\Delta\dot{a}_{ji})}\right| \approx \left|\frac{d_0'(\sigma_0) - 1}{d_0'(\sigma_0)}\right|$$

For $d_0'(\sigma) > 0.5$, $(\Delta\dot{a}_{ji})$ dominates, and we know that the corresponding perturbation $(\delta, 0)$ is unstable in this region where $\sigma < \rho$. For $0 < d_0'(\sigma) < 0.5$, $(\Delta\dot{a}_{ij})$ dominates, and we know that the perturbation (δ, δ) is locally stable because $d_0'(\sigma_0) < 1$ and (δ, δ) lies above the nullcline. This analysis is summarized in Figure 3.2. For $d_0'(\sigma) = 0.5$, the linearized analysis of [5] fails ($\Delta_1 = \Delta_2$ in Figure 3.2), but our analysis easily incorporates higher-order terms:

$$\left|\frac{\dot{a}_{ij}}{\dot{a}_{ji}}\right| \approx \left|\frac{\frac{\delta}{2}d_0''(\sigma_0) + d_0'(\sigma_0) - 1}{\frac{\delta}{2}d_0''(\sigma_0) + d_0'(\sigma_0)}\right|$$

Thus, the collaboration is locally stable for $d_0'(\sigma) = 0.5$ and $d_0''(\sigma) < 0$, while it is unstable for $d_0'(\sigma) = 0.5$ and $d_0''(\sigma) > 0$, which provides a slightly stronger result than that provided in [5].[14] As a final note, we see that for large ρ, we could have another region at small σ in which $d_0'(\sigma) < 0.5$, indicating a second region of local stability to perturbations in a_{ij} for a single bidirectional collaboration. According to Proposition 2.2, however, this level of investment is rarely encountered in stable BCC's.

[14] We could add even higher-order terms as necessary.

3.1.1 Archetypical Network Structures in (σ, d) Space

Having presented the natural form with which to understand network dynamics, we now show two archetypes for network configurations which represent extremes in network dynamics.

Homogeneous Complete BCC We define a homogenous complete BCC as a complete graph in which $a_{ij} = a_{ji} \neq 0$ for all pairs (i,j) within the BCC. By symmetry, $\sigma_{ij} = 2a_{ij}$ and $d_i^{out} = (N-1)a_{ij}$, so the dynamics are confined to the line $d = \frac{N-1}{2}\sigma$ for all pairs (σ_{ij}, d_i^{out}); this is a one-dimensional subspace of the $\frac{N(N+1)}{2}$-dimensional network dynamics $(d_1^{out}, d_2^{out}, ..., \sigma_{12}, \sigma_{13}, ...)$. A few of these lines are shown with the nullcline in Figure 3.3. If we restrict the dynamics to maintain the homogeneous BCC state (nobody actively tries to break symmetry and the network only moves along the line $d = \frac{N-1}{2}\sigma$), then an equilibrium is locally stable when $d_0'(\sigma) < \frac{N-1}{2}$. Thus, the equilibrium can be locally stable to such restricted perturbations even for $\sigma < \rho$, which is not a violation of Proposition 2.2.

Local stability to *arbitrary* perturbations is possible only if $\sigma \geq \rho$ at equilibrium for a homogeneous complete BCC of $N > 2$ (by Proposition 2.2).[15] This observation helps us find two critical values of N which characterize bounds on local stability of homogeneous complete BCC's at equilibrium. For $N \geq N_u$, a homogeneous complete BCC is guaranteed to be unstable to arbitrary perturbations. For $N \leq N_s$, a homogeneous complete BCC at equilibrium can be locally stable or marginally stable to arbitrary perturbations:[16]

$$N_u = \left\lfloor \frac{2}{m\sqrt{r}} + 2 \right\rfloor \tag{3.3a}$$

$$N_s = \begin{cases} \min\left(\left\lfloor \frac{2}{m\sqrt{r}} + 1 \right\rfloor, \lceil 2e_0'(s_s) \rceil \right) & \text{if } r < \frac{9}{16} \\ \left\lfloor \frac{2}{m\sqrt{r}} + 1 \right\rfloor & \text{if } r \geq \frac{9}{16} \end{cases} \tag{3.3b}$$

$$s_s = \frac{5 - \sqrt{9 - 16r}}{8} \tag{3.3c}$$

Here, s, r, and m are the relevant nondimensional forms for σ, τ, and μ. These results are proven in Propositions A.4 and A.5. Intuitively, increasing costs prevent an agent from stably maintaining connections with too many other agents, so a realistic set of parameters m and r gives $N_s \leq c$, where c depends on the problem domain.[17] For $N_s < N < N_u$, there are multiple equilibria, so local stability to arbitrary perturbations depends on the location

[15] The degenerate case of $N = 2$ could be stable anywhere, but the scenario where it reaches a steady-state with $\sigma < \rho$ requires unrealistic choices for r, m.

[16] $N_u > N_s$ but in general $N_u \neq N_s + 1$. For $m = 0.634$ and $r = 0.260$, $N_s = 5$ and $N_u = 8$.

[17] For example, the situation of students trying to solve a homework problem might have $c = 5$ because it is hard to maintain simultaneous real-time collaborations with more than 5 people.

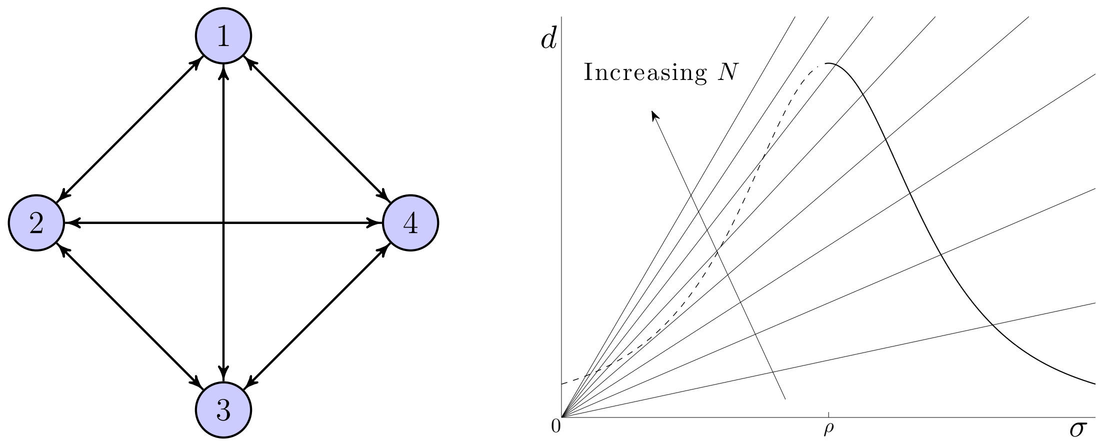

(a) Network structure ($N = 4$) (b) (σ, d) dynamics for $r < \frac{9}{16}$.

Figure 3.3. Homogeneous complete BCC. In (b), lines of $d = \frac{N-1}{2}\sigma$ are shown for some N, implicitly illustrating the pseudo-bifurcation in equilibrium states with variations in the (discrete) parameter N.

of the specific equilibrium on the nullcline ($\sigma < \rho$ at equilibrium implies instability). Viewed in a different light, these results describe bifurcations for the steady states with respect to the parameter N. If N were continuous, there would be (at most two) saddle-node bifurcations at the points where the line $d = \frac{N-1}{2}\sigma$ and the nullcline lie tangent to each other. Because N is discrete, the sequence of lines $d = \frac{N-1}{2}\sigma$ usually jumps over these bifurcation points (as in Figure 3.3) except for very specialized values of m and r.

Symmetric BCC Star We define a BCC star as a BCC in which there is one central node (node 1) with degree $N-1$ and $N-1$ fringe nodes with degree 1. A symmetric BCC star satisfies $a_{1j} = c_1 > 0$, and $a_{j1} = c_2 > 0$, for constants c_1 and c_2 over all fringe nodes j. The dynamics are then constrained to $\sigma_{1j} = d_j^{out} + d_1^{out}/(N-1)$. At equilibrium, $d_j^{out} = d_1^{out}$ (by Proposition 2.1), so the steady state is the intersection of the nullcline with the line $d = \frac{N-1}{N}\sigma$ (see Figure 3.4). Because the slope $\frac{N-1}{N} < 1 \;\forall\; N$, symmetric BCC stars generally have a much wider range of stability in N. We can find the critical value such that when $N \geq N_u^+$ the symmetric BCC star is always unstable to arbitrary perturbations:

$$N_u^+ = \begin{cases} \left\lfloor \frac{m\sqrt{r}}{m\sqrt{r}-1} + 1 \right\rfloor & \text{if } m\sqrt{r} > 1 \\ \infty & \text{if } m\sqrt{r} \leq 1 \end{cases} \tag{3.4}$$

This result is shown in Proposition A.6, which follows the form of Proposition A.4. Intuitively, the star network is easier to manage for each individual, since only the leader (presumably the most capable individual) has to manage simultaneous collaborations. We will revisit this observation in Section 4.1.2 and 4.2 when considering dynamic leader selection.

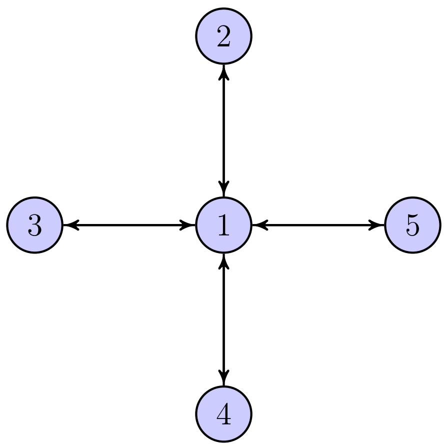

(a) Network structure ($N = 5$)

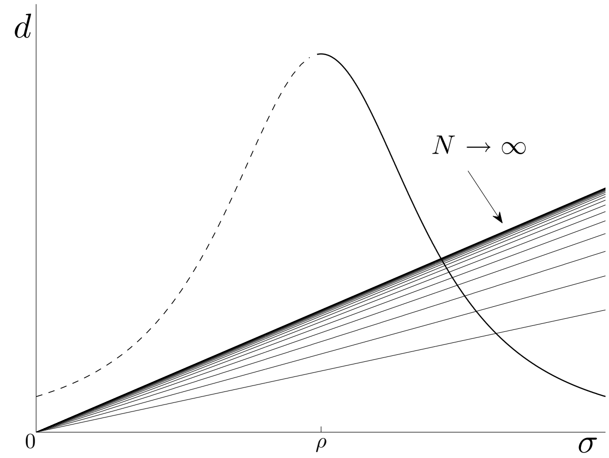

(b) (σ, d) dynamics for $m\sqrt{r} \leq 1$

Figure 3.4. Symmetric BCC star. In (b) lines of $d = \frac{N-1}{N}\sigma$ are shown for many N. These lines approach $d = \sigma$ as $N \to \infty$.

3.2 Hidden Modes of Collaboration

Now we are ready to understand hidden modes of collaboration, linear combinations of a_{ij} that are not captured in the dynamics of network evolution. These modes play important roles in the behavior of both network and state dynamics. In the original space of a_{ij}, the network structure can be represented as the vector $\mathbf{m} = (a_{12}, a_{13}, ...) \in \mathbb{R}_+^u$. The network dynamics are governed by the vector $\mathbf{n} = (d_1^{out}, d_2^{out}, ..., \sigma_{12}, \sigma_{13}, ...) \in \mathbb{R}_+^v$. In general, the dimensions u and v change with the network structure based on the number of a_{ij} or σ_{ij} that are forced to be 0 and which are therefore removed from the dynamics. These two vectors are always linearly related: $\mathbf{n} = T\mathbf{m}$, where $T \in \mathbb{R}^{v \times u}$. We can then define the space of hidden collaborations ($\mathcal{H}$) and the hidden modes of collaboration ($\mathbf{h}$) as follows:

$$\mathcal{H} = \ker T, \ \mathbf{h} \in \mathcal{H} \tag{3.5}$$

3.2.1 Ramifications on Network Dynamics

Marginal Stability The presence of hidden modes that are not captured by the network dynamics yields very important trends which have been mentioned previously. The first is that many network structures can only have marginal stability at best. Consider a complete BCC, in which $u = N(N-1)$ and $v = \frac{N(N+1)}{2}$. Since $\dim \mathcal{H} \geq u - v = \frac{N(N-3)}{2}$, hidden modes are guaranteed for $N > 3$ simply by the rectangular nature of T.[18] When a steady state is

[18] We will shortly show that $\dim \mathcal{H} > 0$ for $N = 3$ as well.

perturbed by $\mathbf{h}$, the network remains stationary, because this perturbation does not move the state away from $d_0(\sigma)$. In other words, $\mathbf{h}$ is literally invisible in (σ, d) space, yielding marginal stability to hidden mode perturbations.

Heterogeneity in Degree-Homogeneous BCC's The second and related ramification on network dynamics is that hidden modes allow heterogeneity in a_{ij} even when the nodes are degree-homogeneous or otherwise indistinguishable from each other. This is evident by noting that a degree-homogeneous BCC at homogeneous equilibrium (such as a homogeneous complete BCC) is completely symmetric. A perturbation by $\mathbf{h}$ necessarily lowers at least one a_{ij} and raises another, because all σ_{ij} and d_i^{out} have to remain the same. The network remains stationary after this perturbation, thereby creating heterogeneity in a_{ij} at steady-state.

Hidden Modes in Network Extremes We saw in the last section that the complete BCC and BCC star represent extremes in network configurations. The former has the most degrees of freedom for any BCC structure of N nodes, while the latter is one of the BCC configurations with the least degrees of freedom.[19] As a consequence, we would expect $\dim \mathcal{H}$ to have extremes over these two structures. We already saw that the lower bound of $\dim \mathcal{H}$ is $\frac{N(N-3)}{2}$ for a complete BCC. For a BCC star, $u = 2(N-1)$ and $v = 2N-1$, so the (σ, d) dynamics have one more dimension than the a_{ij} dynamics. This means that the lower bound of $\dim \mathcal{H}$ for a BCC star is 0. We now determine the number of hidden modes for both network structures by examining T.

Proposition 3.1. *A BCC star has no hidden modes of collaboration.*

Proof. We can reorder the entries in $\mathbf{m}$ and $\mathbf{n}$ any way we please; this will merely permute the rows and columns of T, but it will not change $\operatorname{rank} T$. Let us write the first $N-1$ entries of $\mathbf{m}$ as a_{1j} and the last $N-1$ entries as a_{j1}. The first entry of $\mathbf{n}$ is d_1^{out}, the next $N-1$ entries are σ_{1j}, and the last $N-1$ entries are d_j^{out}. Then, T can be written as follows:

$$T = \left(\begin{array}{c|c} \mathbf{1}_{N-1}^T & \mathbf{0}_{N-1}^T \\ \hline I_{N-1} & I_{N-1} \\ \hline 0_{N-1,N-1} & I_{N-1} \end{array} \right) = \left(\begin{array}{c} \mathbf{g}^T \\ \hline P \end{array} \right)$$

where $0_{i,j}$ is the zero matrix of size $i \times j$, $\mathbf{g} \in \mathbb{R}^{2N-2}$ and $P \in \mathbb{R}^{(2N-2)\times(2N-2)}$. The triangular form of P implies that $\operatorname{rank} P = 2(N-1)$, and it is clear that $\mathbf{g}^T$ is the sum of the first $N-1$ rows of P subtracted by the sum of the last $N-1$ rows. Therefore $\operatorname{rank} T = 2(N-1)$, and $\dim \mathcal{H} = u - \operatorname{rank} T = 0$. □

[19] A path also has the same number of degrees of freedom.

Proposition 3.2. *A complete BCC has* $\binom{N-1}{2}$ *hidden modes of collaboration.*

Proof. To make notation clear, we temporarily add subscripts $*_N$ to denote variables associated with a system of N nodes. Then, we can relate T_N to T_{N+1}. We append $a_{j,N+1}$ followed by $a_{N+1,j}$ to $\mathbf{m}$ (where $j \in \{1, 2, ..., N\}$). Also, we append the entry d^{out}_{N+1} followed by $\sigma_{j,N+1}$ to $\mathbf{n}$ (where $j \in \{1, 2, ..., N\}$). Then, T_{N+1} is structured in the following way (where the base case of T_1 is the empty matrix):

$$T_{N+1} = \left(\begin{array}{c|c} T_N & X_N \\ \hline \mathbf{0}^T_{N(N-1)} & \mathbf{k}^T \\ \hline 0_{N,N(N-1)} & Y_N \end{array}\right)$$

$$X_N \in \mathbb{R}^{v_N \times 2N}, \quad \mathbf{k} = \left(\begin{array}{c} \mathbf{0}_N \\ \hline \mathbf{1}_N \end{array}\right), \quad Y_N = \left(\begin{array}{c|c} I_N & I_N \end{array}\right)$$

where X_N has N nonzero entries: there is a 1 in each row/column pair corresponding to the mapping from $a_{N+1,j}$ to d^{out}_j ($j \in \{1, 2, ..., N\}$). We can write d^{out}_{N+1} as a linear combination of the other variables in $\mathbf{n}$:

$$d^{out}_{N+1} = \sum_{i,j} \sigma_{ij} \; - \; \sum_{i=1}^{N} d^{out}_N$$

Thus, we can eliminate the corresponding row $\left(\mathbf{0}^T_{N(N-1)} \mid \mathbf{k}^T\right)$ when computing the rank of T_{N+1}, and we now know $\operatorname{rank} T_N \leq v_N - 1$. The N rows in the submatrix $\left(0_{N,N(N-1)} \mid Y_N\right)$ are linearly independent with each other by inspection of the form of Y_N. Also, every row in $\left(0_{N,N(N-1)} \mid Y_N\right)$ has a nonzero element in the last N columns, whereas the last N columns of X_N are all zeros. Therefore the rows in $\left(0_{N,N(N-1)} \mid Y_N\right)$ are also linearly independent with the rows of the submatrix $(T_N \mid X_N)$. Applying this argument recursively back to $N = 1$, we can see that all rows in T_{N+1} corresponding to σ_{ij} (where $i \in \{1, 2, ..., N\}$ and $j \in \{2, 3, ..., N + 1\}$) are linearly independent with each other, and they are linearly independent with all rows corresponding to d^{out}_i (where $i \in \{1, 2, ..., N\}$). This also implies

$$\operatorname{rank} T_{N+1} = \operatorname{rank} (T_N \mid X_N) + \operatorname{rank} \left(0_{N,N(N-1)} \mid Y_N\right) = \operatorname{rank} (T_N \mid X_N) + N.$$

We can permute the rows in $(T_N \mid X_N)$ such that the first N rows of X_N look like the matrix $(I_N|0_{N,N})$. Then these rows, which correspond to d^{out}_i (where $i \in \{1, 2, ..., N\}$), are linearly independent with each other. So we have just shown that all of the rows in $(T_N \mid X_N)$, which are composed of d^{out}_k and σ_{pq} (where $k \in \{1, 2, ..., N\}$, $p \in \{1, 2, ..., N-1\}$, and $q \in \{2, 3, ..., N\}$), are linearly independent: $\operatorname{rank} (T_N \mid X_N) = v_N$. We now have the

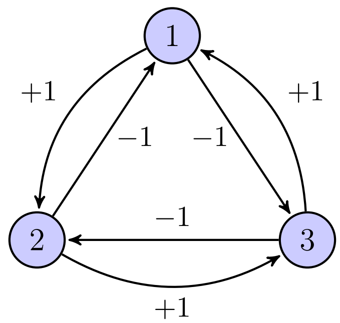

Figure 3.5. A symmetric zero-sum 3-cycle: $\mathbf{h}^{c3}_{1,2,3}$. Perturbing a system by $c\mathbf{h}^{c3}_{1,2,3}$ implies increasing the magnitudes of the edges a_{12}, a_{23}, and a_{31} by c while reducing the magnitudes of the edges a_{13}, a_{32}, and a_{21} by c.

difference equation: $\operatorname{rank} T_{N+1} = v_N + N = v_{N+1} - 1$, so $\operatorname{rank} T_N = v_N - 1 \; \forall \; N > 0$. Thus, $\dim \mathcal{H} = u - \operatorname{rank} T = u - v + 1 = \frac{(N-1)(N-2)}{2} = \binom{N-1}{2}$. $\square$

We can summarize these two propositions as follows. The dimension of $\mathcal{H}$ is a metric for what we qualitatively introduced as network extremes: a BCC star has no hidden modes, whereas a complete BCC has a hidden space that grows $\sim \frac{N^2}{2}$, i.e. nearly half of the a_{ij} modes are hidden for large N. More importantly, the space of hidden modes for a complete BCC with N nodes contains the space of hidden modes for any other graph of N nodes, i.e. we can have at most $\binom{N-1}{2}$ hidden modes for any graph of N nodes. This is because any graph can be created by just removing edges from a complete graph with the same number of nodes. In other words, if a hidden mode exists on any graph of N nodes, it must exist on the complete BCC of N nodes.

A Metric for Network Redundancy We have established that $\dim \mathcal{H}$ is a metric with which we may quantitatively differentiate between different types of graphs, but we have not yet defined what this metric measures. Intuitively, we expect the hidden modes to represent redundancy, and this is precisely what we discover by finding a basis for $\mathcal{H}$. The complete BCC with $N = 3$ has one hidden mode. If we set the order as in Proposition 3.2, then $\mathbf{m} = (a_{12}, a_{21}, a_{13}, a_{23}, a_{31}, a_{32})^T$, and $\mathbf{n} = (d_1^{out}, d_2^{out}, \sigma_{12}, d_3^{out}, \sigma_{13}, \sigma_{23})^T$. The (unnormalized) fundamental hidden mode is:

$$\mathbf{h}^{c3}_{1,2,3} = (1, -1, -1, 1, 1, -1)^T$$

We call this structure a symmetric zero-sum 3-cycle ($\mathbf{h}^{c3}$), because it is composed of 3 nodes (labeled counterclockwise in subscripts) joined in a symmetric cycle with a net sum of 0 over the edge weights. The vector $\mathbf{h}^{c3}_{1,2,3}$ is shown Figure 3.5.

In Proposition A.7 we prove that for any graph, $\mathcal{H}$ is spanned by symmetric zero-sum 3-cycles. In general, many of these $\binom{N}{3}$ vectors are linearly dependent: $\mathbf{h}^{c3}_{1,2,3} - \mathbf{h}^{c3}_{1,2,4} + \mathbf{h}^{c3}_{1,3,4} = \mathbf{h}^{c3}_{2,3,4}$. Also, $\mathbf{h}^{c3}_{1,2,3} - \mathbf{h}^{c3}_{2,3,4} = \mathbf{h}^{c4}_{1,2,4,3}$, which is a symmetric zero-sum 4-cycle. In this way, linear combinations of symmetric zero-sum 3-cycles produce all types of hidden modes, including larger symmetric zero-sum cycles. Now we can relate hidden modes to network structure: a network has no hidden modes iff there is no way to make a symmetric zero-sum k-cycle, where $3 \leq k \leq N$. Furthermore, the number of hidden modes represents the number of linearly independent $\mathbf{h}^{ck}$ that exist for a graph. A symmetric zero-sum k-cycle formally characterizes redundancy because it represents the ability for information output to flow back to the source as input without retracing an edge. It represents *indirect* feedback in a network. In contrast, *direct* feedback is represented by any bidirectional link.

3.2.2 Ramifications on State Dynamics

In and of itself, redundancy (i.e. indirect feedback) may be beneficial or detrimental to multi-agent systems. In the context of consensus dynamics, we could expect redundancy to be detrimental to the speed of convergence, because agents waste time and resources receiving information that they already know. The effects on robustness are less obvious, since "turning the knobs" on the indirect feedback terms (i.e. changing the magnitudes of the hidden mode vectors) may increase robustness up to a certain point after which the feedback overcompensates for perturbations. These intuitive notions are easily cast into a more formal framework: the H_2-norm (abbreviated as H_2) and the speed of convergence (abbreviated as λ_{min}) change as we perturb a system along a hidden mode $\mathbf{h}$. Analyzing their variations is similar to looking at the root locus plot of a transfer function; instead of tracking the change in poles of a transfer function as we change a parameter, we can track the changes in the aforementioned characteristics of state dynamics as we perturb A with $c\mathbf{h}$, where c is a real number bounded by the fact that $a_{ij} \geq 0$. We show two examples of these variations below.

When presenting such results, it is convenient to compare the changes we observe to a baseline value, one that is independent of the hidden mode perturbations. The homogeneous complete BCC has a very special structure; we now show that the state dynamics for a homogeneous complete BCC (henceforth called H_2^* and λ^*_{min}) are unaffected by hidden mode perturbations.

Proposition 3.3. *H_2^* and λ^*_{min}, the H_2-norm and convergence speed for a homogeneous complete BCC, are invariant under hidden mode perturbations.*

Proof. Consider perturbing the adjacency matrix A of a homogeneous complete BCC by a linear combination of all $\binom{N-1}{2}$ linearly independent hidden modes: $\sum_k c_k \mathbf{h}^{c3}_{m_k,n_k,o_k}$, where

m_k, n_k, and o_k are the three nodes of the $k-$th symmetric zero-sum 3-cycle. Because $d_i^{out} = d_j^{out} = d^{out}$ over all pairs of nodes (i, j) (all weighted outdegrees are the same), the perturbed normalized Laplacian ($\hat{L}_p$) can be written in the following way:

$$\hat{L}_p = \hat{L} - C \ , \ C = \frac{1}{d^{out}} \sum_k c_k \mathbf{h}^{c3}_{m_k, n_k, o_k}$$

where $\hat{L}$ is the unperturbed normalized Laplacian. The unperturbed normalized Laplacian is symmetric and therefore normal and diagonalizable, and it looks like the following:

$$\hat{L} = \frac{1}{N-1}(NI_N - \mathbf{1}_N \mathbf{1}_N^T) = \frac{N}{N-1}\Pi$$

The matrix C is skew-symmetric, so it is also normal and diagonalizable. Because $\hat{L}$ is a projection matrix onto the subspace orthogonal to $\mathbf{1}_N$ and because C is already orthogonal to $\mathbf{1}_N$ by construction, we must have $\hat{L}C = C\hat{L} = C$. Therefore, the two matrices $\hat{L}$ and C can be simultaneously diagonalized, whereby the eigenvalues of $\hat{L}_p$ are just the eigenvalues of $\hat{L}$ subtracted by those of C. Since C is skew-symmetric, it has imaginary eigenvalues. Thus, λ^*_{min} does not change under hidden mode perturbations. The matrix $\hat{L}_p$ is also normal because it is the sum of normal matrices which commute. Then, by Proposition 1 of [16], we can write the H_2-norm of $\hat{L}_p$ as:

$$H_2^* = \left(\sum_{i=2}^{N} \frac{1}{2 \operatorname{Re} \lambda_i} \right)^{\frac{1}{2}}$$

The real parts of the eigenvalues for $\hat{L}$ and $\hat{L}_p$ are the same, so the H_2-norm is also the same. Thus, H_2^* and λ^*_{min} are invariant under hidden mode perturbations. By symmetry, all nonzero eigenvalues are the same for $\hat{L}$ of a homogeneous complete BCC. Also, $\operatorname{tr} \hat{L} = N$, so the values of λ^*_{min} and H_2^* are:

$$\lambda^*_{min} = \frac{N}{N-1} \tag{3.6a}$$

$$H_2^* = \frac{N-1}{\sqrt{2N}} \tag{3.6b}$$

□

Consider a non-homogeneous complete BCC of 3 nodes, which has just one hidden mode: $\mathbf{h}^{c3}_{1,2,3}$. We create a random matrix where a_{ij} are chosen from the Gaussian distribution

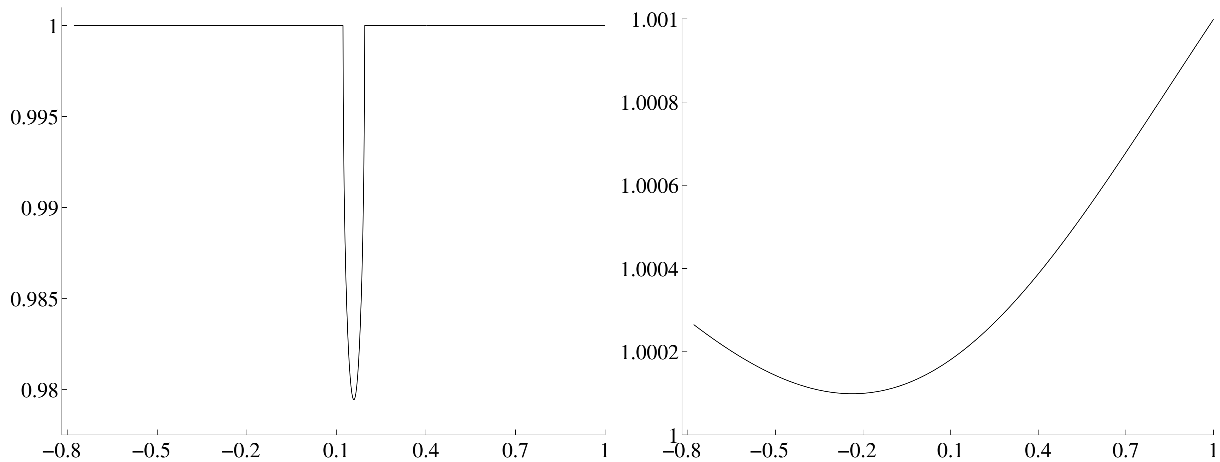

(a) $\lambda_{min}/\lambda^*_{min}$ vs. c (b) H_2/H_2^* vs. c

Figure 3.6. The effect of perturbing a complete BCC ($N = 3$) with the hidden mode $\mathbf{h}^{c3}_{1,2,3}$.

$\mathcal{N}(1, 0.1^2)$. Then the perturbed A matrix may look like the following:

$$A = \begin{pmatrix} 0 & 0.78 & 1.00 \\ 1.31 & 0 & 0.85 \\ 0.95 & 1.23 & 0 \end{pmatrix} + c \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \quad -0.78 \leq c \leq 1.00$$

Figure 3.6 shows H_2/H_2^* and $\lambda_{min}/\lambda^*_{min}$ with variations in c. The shape of the speed of convergence plot always looks like a horizontally shifted version of Figure 3.6(a) for complete graphs of $N = 3$ (non-complete graphs do not even have hidden modes for $N = 3$). There are only 3 eigenvalues for $\hat{L}$ and one of them is forced to be zero by construction, so only two eigenvalues change. The homogeneous complete graph has all nonzero eigenvalues equal to λ^*_{min}, so the nonzero eigenvalues of any connected graph of 3 nodes must be symmetrically placed about λ^*_{min}. Therefore, $\lambda_{min} \leq \lambda^*_{min}$, as shown in Figure 3.6. The value of c which corresponds to the minimum speed of convergence ($c \approx 0.19$) is the point at which $\hat{L}$ has eigenvalues "as real as possible," or the skew-symmetric component of $\hat{L}$ is minimized:

$$\arg\min_c \frac{\lambda_{min}}{\lambda^*_{min}} = \arg\min_c \left\| \frac{\hat{L} - \hat{L}^T}{2} \right\|_2$$

The general behavior of the H_2-norm plot is more complicated. The global behavior of the plot (if we let $c \in \mathbb{R}$ irrespective of the non-physicality for negative elements in A) can have a global maximum (but no other extrema), a global minimum and a global maximum, or only a global minimum (but no other extrema) as shown in Figure 3.6(b); it always has a horizontal asymptote as well. This non-monotonic behavior is due to the fact that

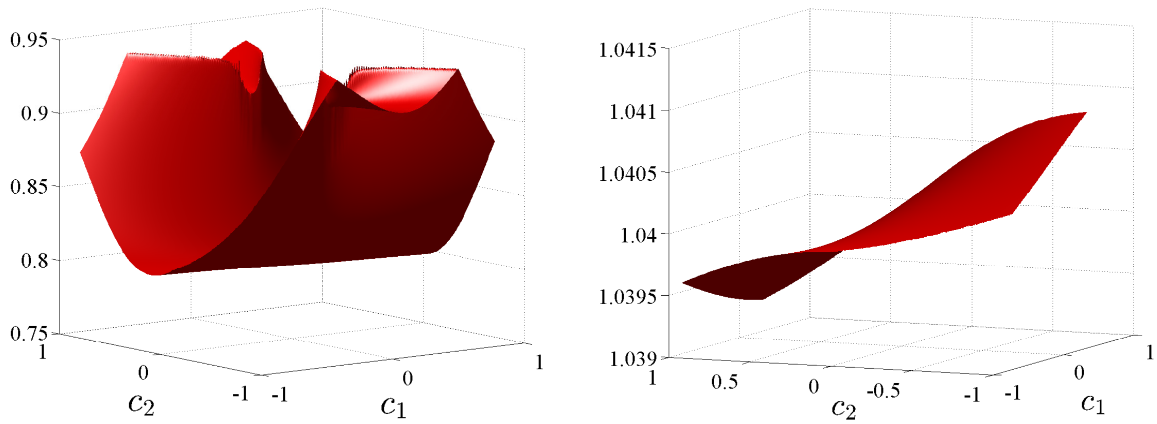

(a) $\lambda_{min}/\lambda^*_{min}$ vs. c_1 and c_2

(b) H_2/H_2^* vs. c_1 and c_2

Figure 3.7. The effect of perturbing a complete BCC ($N = 4$) with the linear combination of hidden modes $c_1\mathbf{h}^{c3}_{1,2,3} + c_2\mathbf{h}^{c3}_{1,2,4}$.

the behavior of the H_2-norm depends on the solution of a Lyapunov equation of 2×2 matrices, which yields at most fourth-order variations for elements of the solution matrix. The asymptote is expected because the system looks very similar to $\pm\mathbf{h}^{c3}_{1,2,3}$ at large $|c|$. Despite this complicated behavior, the fact that H_2 varies non-monotonically with c verifies our intuition about redundancy and robustness (at least for $N = 3$).

Moving to $N > 3$ yields even more complex behavior for the response of state dynamics to hidden modes. We append the 3×3 matrix above with the column $(1.09, 0.71, 1.14, 0.98)^T$ and the row $(0.83, 0.99, 1.03, 0.98)$, and perturb this 4-node system with the linear combination $c_1\mathbf{h}^{c3}_{1,2,3} + c_2\mathbf{h}^{c4}_{1,2,4}$ in Figure 3.7.[20] The speed of convergence no longer has the simple shape as in the case of Figure 3.6: now it is the minimum of the superposition of two convex surfaces. Nevertheless, some observations from $N = 3$ still apply: $\lambda_{min} \leq \lambda^*_{min}$ and $\min_{c_1,c_2} \lambda_{min}$ still corresponds to the point which minimizes the skew-symmetry of $\hat{L}$.[21] The magnitudes of deviation from λ^*_{min} can be greater as we increase the degrees of freedom; each hidden mode affects a pair of eigenvalues, so the hidden modes can couple on a single eigenvalue. The variation of H_2 remains non-monotonic and in this case shows saddle-like behavior over the parameter domain. The magnitude of H_2 variations also increases due to the increased degrees of freedom. Although we cannot definitively comment on the shapes of their variations under hidden mode perturbations as N increases even further, we can always bound these plots by the following results: $\lambda_{min} \leq \lambda^*_{min}$ and $H_2 \geq H_2^*$ for any connected graph (see Theorems A.8 and A.9).

[20] Note that the constraints on c_2 are in general dependent on c_1.

[21] Although it is hard to see, the curve in Figure 3.7 has a well-defined minimum at $(c_1, c_2) \approx (0.33, 0.11)$.

Our analysis verifies intuitive notions about the effects of redundancy on the speed of convergence and robustness of a network with respect to consensus. The fact that redundancy affects state dynamics but not network dynamics is a fundamental weakness in the utility model. Before presenting ways in which to improve upon the utility model, we analyze a reduced-order system to display another weakness: the ability of a connected graph to break and become disconnected.

3.3 Reduced Dynamics of Complete BCC's

Realistic systems involve large N and such systems become extremely complicated to analyze analytically due to the nonlinear and highly coupled nature of the dynamics. Therefore, it is important to be able to analyze a reduced-order system, a model that necessarily throws away some of the complexity of the original model but renders the detailed analysis of arbitrarily large systems tractable and approachable.

Rather than constrain the dynamics to be linear (which would grossly change the behavior of each node), we leave the utility maximization protocol (Equation 2.3b) intact. Instead, we constrain the network's initial conditions. We begin with a complete BCC because all other graph types are contained within this network, and it has the greatest possibilities for network evolution. Also, we prescribe that $N > 3$ since we are concerned with large-scale networks, and we add symmetry to the initial condition by enforcing that the network is *almost* homogeneous. Namely, we perturb just one edge (a_{12}) of a homogeneous complete BCC. This results in a 7th order system for any $N > 3$, depicted graphically in Figure 3.8. Here, the looped edge a_{33} represents the values a_{ij}, where $i \neq j$ and $i, j > 2$. Because these nodes are neither receivers nor producers of the perturbation a_{12}, all such nodes can be grouped into a single node by symmetry. This simplified model allows for more tractable analysis without removing too much of the initial richness of the network dynamics. This model has an associated 7th order system in (σ, d) space due to the addition of σ_{33}. The only hidden mode is similar to $\mathbf{h}^{c3}_{1,2,3}$ but edges a_{12} and a_{21} are scaled by a factor of $N-2$ to account for the fact that node 3 actually represents $N-2$ nodes. Since the edges between members of group 3 are identical in the reduced system (they are all represented by a_{33}), there are no hidden modes purely amongst members of this group.

3.3.1 Linearized Dynamics Near Equilibrium

The linearized dynamics near an equilibrium provide intuition for the behavior of large-scale systems, particularly for $N \geq N_u$. As before, we have the vector $\mathbf{m}$ of a_{ij} (now including a_{33}), and we define the state at equilibrium for a homogeneous complete BCC as $\mathbf{m}_0$. The

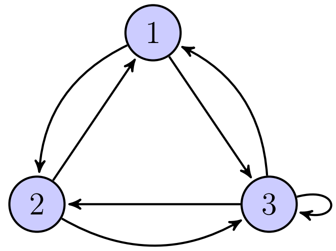

Figure 3.8. Reduced 7th order system

state away from equilibrium is $\boldsymbol{\eta} = \mathbf{m} - \mathbf{m}_0$. Rewriting Equation 2.3b as $\dot{\mathbf{m}} = f(\mathbf{m})$, the linearized dynamics can be written as:

$$\dot{\boldsymbol{\eta}} \approx \left(\frac{\partial f(\mathbf{m})}{\partial \mathbf{m}} \right) \Bigg|_{\mathbf{m}_0} \boldsymbol{\eta} \tag{3.7}$$

where $\frac{\partial f(\mathbf{m})}{\partial \mathbf{m}}$ is a shorthand for the 7×7 Jacobian matrix. These dynamics can also be written in (σ, d) space with an analogous 7×7 Jacobian associated with $\dot{\mathbf{n}}$, where $\mathbf{n} = T\mathbf{m}$ as before. However, because T is singular, the standard change-of-basis transformation cannot be used and the transformation must be done manually prior to linearization (both Jacobians are shown in Appendix B). The most important observations for the linearized system are the eigenvalues (Table 3.1) and associated eigenvectors of perturbations, which evolve into the stable/unstable manifolds for the nonlinear dynamics.[22] Here, the value $d_0'(\sigma_0) = B''(\sigma_0)/2\mu$ denotes the slope of the nullcline at the value σ_0 associated with an equilibrium for a homogeneous complete BCC.

Table 3.1. Eigenvalues for Linearized Reduced Dynamics of a Complete BCC

λ_0	0
λ_1	$4\mu d_0'(\sigma_0)$
λ_2	$4\mu \left(d_0'(\sigma_0) - \frac{N-1}{2} \right)$
λ_3	$2\mu \left(d_0'(\sigma_0) - \frac{N-1}{2} + \sqrt{\left(d_0'(\sigma_0) - \frac{N-1}{2} \right)^2 + N d_0'(\sigma_0)} \right)$
λ_4	$2\mu \left(d_0'(\sigma_0) - \frac{N-1}{2} - \sqrt{\left(d_0'(\sigma_0) - \frac{N-1}{2} \right)^2 + N d_0'(\sigma_0)} \right)$

All dynamics near an equilibrium are governed by the timescale $\frac{1}{\mu}$, which differs from ρ^2 by the (usually $O(1)$) factor m. Because near-equilibrium dynamics are already centered about $\sigma = \rho$, we expect a timescale other than ρ^2. As mentioned in Section 2.1.1, the distance of the equilibrium points from ρ is governed by μ, so it is natural for this to be the relevant scaling factor for the perturbed dynamics.

[22] The eigenvalues λ_3 and λ_4 each have an algebraic and geometric multiplicity of 2.

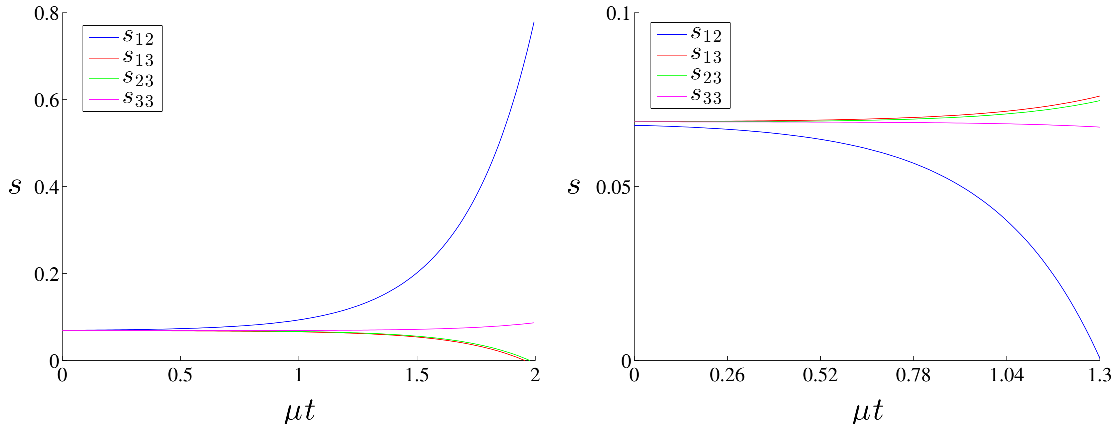

(a) $\delta b_{12} = \delta a_{12}/\rho = 0.001$ (b) $\delta b_{12} = \delta a_{12}/\rho = -0.001$

Figure 3.9. Linearized reduced dynamics for $\rho = 0.65$, $\tau = 0.11$, $\mu = 1.5$, and $N = 11 > N_u = 8$. Both σ and time are shown in dimensionless form. The changes in $e_i^{out} = d_i^{out}/\rho$ are $O(10^{-3})$ over this timescale.

Also, we readily observe many of the same characteristic values noted for the stability of the original, non-reduced system when finding the values for N_u and N_s: λ_0 corresponds to the hidden mode eigenvector, $\lambda_1, \lambda_3 > 0$ when $\sigma_0 < \rho$ (i.e. $N > N_s$), and the condition that $\lambda_2 > 0$ is only possible for $N_s < N < N_u$ (the associated eigenvector is a uniform growth of all a_{ij}). The value for λ_4 is guaranteed to be $< 0 \; \forall \; N > 0$. For $N \geq N_u$, λ_1 and λ_3 are the unstable eigenvalues, and λ_1 dominates behavior ($\lambda_1 > \lambda_3$). Interestingly, the associated eigenvector for λ_1 is the following:

$$\left(0, 0, 0, \frac{(N-2)(N-3)}{2}, \frac{3-N}{2}, \frac{3-N}{2}, 1\right)^T$$

with the variables ordered as $\mathbf{n} = (d_1^{out}, d_2^{out}, d_3^{out}, \sigma_{12}, \sigma_{13}, \sigma_{23}, \sigma_{33})$. Thus, the dominant behavior of a perturbation from an unstable homogeneous equilibrium is a reallocation of investments without any change in individual nodal costs. Namely, σ_{12} and σ_{33} move in a direction opposing σ_{13} and σ_{23}, and the magnitude of σ_{12} grows/dies much faster than any of the other variables for large N. The linearized dynamics for positive and negative δa_{12} (the initial perturbation) are shown in Figure 3.9 for a specific case with $N > N_u$. The dominance of the λ_1 eigenmode is clearly illustrated. In the case where $\delta a_{12} > 0$, we see that σ_{13} and σ_{23} vanish. This results in a disconnected network, which has enormous ramifications regarding state dynamics: consensus is virtually impossible in this case since nodes 1 and 2 isolate themselves from everyone else. More formally, as the system approaches disconnection, $\lambda_{min} \to 0$ and $H_2 \to \infty$. We now consider the fully nonlinear dynamics to see if graph disconnection actually occurs.

3.3.2 Extension to Nonlinear Perturbations

The nonlinear simulations of the cases introduced in Figure 3.9 are shown in Figures 3.10 and 3.11.[23] We are mainly concerned with the steady-state structure of the reduced system. As shown in the figures, the case with a positive initial perturbation results in a disconnected steady-state graph. The latter case with negative initial perturbation results in a connected graph that is no longer a BCC due to the vanishing of link a_{32}.[24] Simulations have shown that these steady-state structures result for any reduced system that does not necessarily start close to equilibrium, so long as the perturbation from homogeneity (i.e. δa_{12}) is small. This observation results from the following argument: the nonlinearity in the dynamics comes from $B'(\sigma)$, which is independent of d. Thus, a vertical perturbation in (σ, d) space only affects the dynamics linearly. A change in σ as we move along the nullcline merely changes the values of $d'(\sigma_0)$ in the eigenvalues for the linearized dynamics. Therefore, we can decompose an initial condition that begins far away from the nullcline but still very close to $d = \frac{N-1}{2}\sigma$ (i.e. small δa_{12}) as a shift along the nullcline followed by a vertical shift in d. In this way, we find that the dynamics move along the eigenvector associated with λ_2 (the fastest negative pole), which implies a uniform decrease in a_{ij} until the system gets close to the nullcline. Taking this point to be the effective initial condition, we can approximate further time evolution by the linearized system; δa_{12} is further than the other δa_{ij} from the nullcline and the line $d = \frac{N-1}{2}\sigma$ after the system moved along the λ_2 eigenmode. Again, this result requires δa_{12} from the line $d = \frac{N-1}{2}\sigma$ is small. Large perturbations δa_{12} from the line $d = \frac{N-1}{2}\sigma$ result in widely different behaviors depending on the specific initial condition.

We can also argue that the linearization of the reduced system approximates the real system well enough that we should expect graph breakage for all reduced systems with $N > N_u$ and $\delta a_{12} > 0$ (but $\delta a_{12} << a_{12}$). For $s < 1$, the only nonlinear terms that could appreciably come into play for the dynamics away from equilibrium are the terms proportional to $B'''(\sigma)$ (the first neglected terms from the linearizations). Consider a generic perturbation δ. The first term neglected from the linearization compares to the corresponding terms in the Jacobian in the following way:

$$\frac{\frac{1}{2}B'''(s)\delta^2}{B''(s)\delta} = \frac{\delta}{2}\left[\frac{r - 4z^2}{z(r + z^2)}\right] \tag{3.8}$$
$$z = s - 1$$

The neglected term becomes increasingly important as we move towards $z = 0$ (i.e. $\sigma = \rho$).

[23] We use a fourth-order Runge-Kutta method to solve the system of ODE's.

[24] Note that (σ_{23}, e_2) ends at a location neither on the d axis nor on the nullcline. Stationarity for a vanished link can occur anywhere in (σ, d) space.

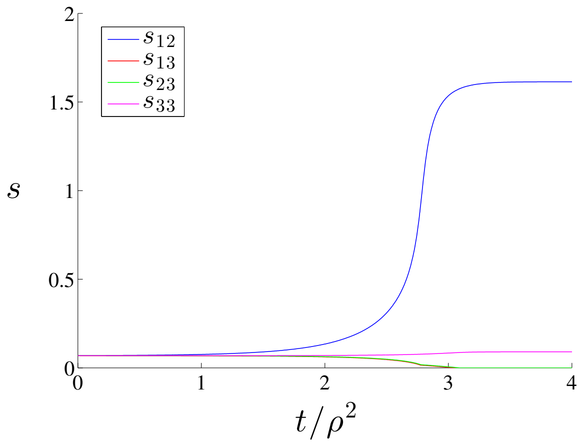

(a) Dimensionless σ vs. time

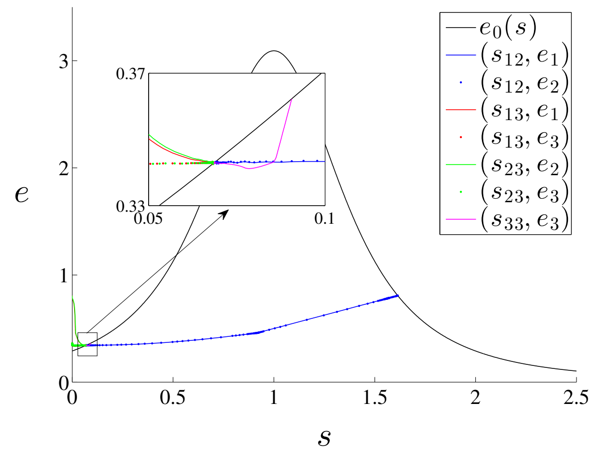

(b) Dimensionless (σ, d) dynamics

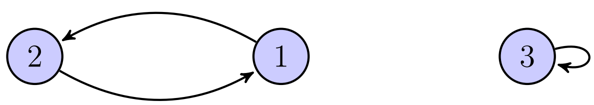

(c) Steady-state structure

Figure 3.10. Nonlinear reduced dynamics for $\delta b_{12} = 0.001$, $\rho = 0.65$, $\tau = 0.11$, $\mu = 1.5$, and $N = 11 > N_u = 8$. The steady state graph is disconnected, as predicted by linearized dynamics.

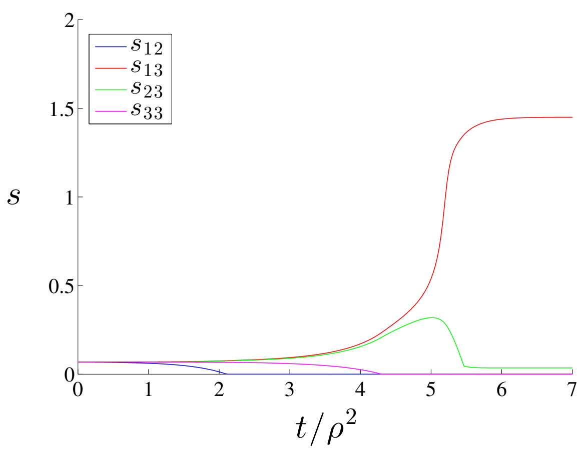

(a) Dimensionless σ vs. time

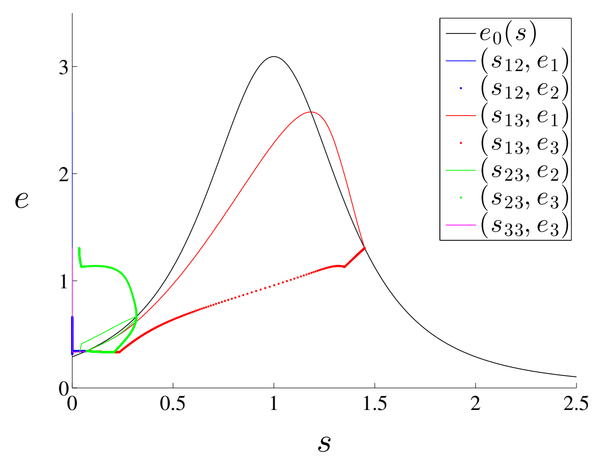

(b) Dimensionless (σ, d) dynamics

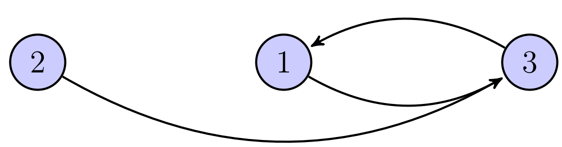

(c) Steady-state structure

Figure 3.11. Nonlinear reduced dynamics for $\delta b_{12} = -0.001$, $\rho = 0.65$, $\tau = 0.11$, $\mu = 1.5$, and $N = 11 > N_u = 8$. The steady-state graph is connected, but 4 of the 7 edges have vanished. Node 2 is a fringe node looking at group 3 through a unidirectional link.

As $z \to 0^{\pm}$, the ratio in Equation 3.8 approaches $\pm\infty$. As we move away from $z = 0$ and hit $z = \pm\frac{\sqrt{r}}{2}$ (the points of inflection for the nullcline), the ratio drops to zero. Moving even further, the ratio has a maximum (minimum) which goes as $O(\delta/\sqrt{r})$ at a location $|z| = O(\sqrt{r})$ for $z < 0$ ($z > 0$), after which the ratio approaches zero for larger $|z|$. Considering the case of $N > N_u$, we must have $-1 < z < 0$ at the initial near-equilibrium condition, and this point moves closer to $z = -1$ as N increases. Because σ_{13} and σ_{23} dynamics move towards more negative z from the initial equilibrium point for $\delta a_{12} > 0$, the nonlinear terms do not appreciably affect their dynamics. The nonlinear terms definitely affect σ_{12} since the point of inflection for the time evolution of σ_{12} occurs near $\sigma = \rho$. However, because the d variables change negligibly as long as z is not close to 0, the points (s_{12}, e_1) and (s_{12}, e_2) are likely to lie underneath the nullcline when $z \to 0^-$ (i.e. $s \to 1^-$). Therefore, the nonlinear term only slows down the motion of σ_{12} but cannot reverse it. To see this more formally, we can write:

$$\dot{\sigma}_{ij} = 2B'(\sigma_{ij}) - 2\mu(d_i^{out} + d_j^{out}) = 4\mu\left(d_0 - \frac{d_i^{out} + d_j^{out}}{2}\right)$$

Thus, σ_{12} is guaranteed to increase as long as the average of the costs for nodes 1 and 2 lies below the nullcline. At $z > 0$, the nonlinear term only helps push σ_{12} towards the direction that the linearization already wants to move. In this way, the linearization approximates the (reduced) nonlinear behavior very well. By the time the nonlinear terms come into play, it is already too late to drastically alter trajectories decided by the linearization.

Importantly, this analysis employs properties specific to the reduced system and its linearization, and it does not always hold for arbitrary initial conditions or networks. Nevertheless, it helps us justify another serious drawback to the utility model: the model allows for graph breakage not just in unique cases, but over a whole category of graphs and initial conditions.

Main Deficiency in the Utility Model

Our analysis of the network and state dynamics has revealed two important problems regarding the utility model for network evolution. The first concerns hidden modes: redundancy is not adequately captured by network dynamics, which results in marginal stability for many classes of networks and the inability for the network to respond to certain network perturbations despite changes in state dynamics. When considering this model in the framework of consensus dynamics, such behavior is simply unacceptable. The second, and perhaps more serious, problem is that the utility maximization protocol allows for graph breakage; this is a major concern since graph breakage renders consensus virtually impossible.

Both of these problems result from the fact that the network dynamics are not coupled to the state dynamics. Whereas the state dynamics depend on the network through $\hat{L}$, the utility protocol currently makes no reference to the state dynamics when considering benefits or costs of collaboration. We now consider ways in which to introduce such coupling so that the network dynamics incorporate information about the systems's state, thereby addressing the problems with redundancy and graph breakage. Importantly, we still require that this coupling remain local such that the overall consensus protocol remains decentralized.

4 Improvements of the Utility Model for Noisy Consensus Dynamics

Our main research goal is to develop a decentralized protocol for noisy consensus dynamics using a utility maximization approach. So far, we have introduced a utility model that captures many realistic aspects of sociological network behavior and therefore appears promising as an engineering design. Upon further analysis, we have shown that the utility model currently lacks coupling with state dynamics, which leads to poor performance with respect to developing a robust and speedy consensus. We now present methods to sophisticate the utility model.

4.1 Perceived Intelligence

Consider again the situation of students trying to solve a homework problem, which we used to motivate the utility model in Section 2.1. An efficient individual will tend to spend more resources collaborating with a neighbor who appears to be smarter or at least more confident in his/her own work. Without an a priori sense for how smart or confident each student actually may be, intelligence can only be observed through the measure of how close to consensus a student is with respect to his/her neighbors: an inherently smart student will tend to develop consensus with his/her neighbors quickly, and, conversely, those who are close to consensus with their neighbors will quickly develop confidence in their work. We use this intuition to motivate the "perceived intelligence" factor (f) with which to weight the benefit of a collaboration.

We define the set of neighbors for node k ($\mathcal{N}_k$) as the nodes l which satisfy $a_{kl} > 0$ or $a_{lk} > 0$ in the initial condition of the network. This specification captures the fact that even if node k no longer chooses to collaborate with node l, it can still observe l's intelligence just by being adjacent. In other words, neighbors are adjacent members in the undirected version of the initial network state. We can then define the unweighted degree of node k as $\deg(k)$, which counts the elements in $\mathcal{N}_k$ and is not to be confused with the weighted degrees d_k^{out} or d_k^{in}. The variance of the state of node k with respect to its neighbors can be written as:

$$v_k = \sum_{l \in \mathcal{N}_k} \frac{(x_k - x_l)^2}{\deg(k)}$$

Then, the absolute intelligence of node k goes as $O(1/v_k)$.[25] Node $i \in \mathcal{N}_k$ observes the

[25] The fact that this parameter diverges at consensus will be dealt with shortly.

absolute intelligence of node k and all of the other nodes in $\mathcal{N}_i$. Node i then evaluates this value against the average intelligence in $\mathcal{N}_i$, resulting in the perceived intelligence factor f_{ik}:

$$f_{ik} = \frac{\frac{1}{v_k}}{\frac{1}{\deg(i)}\sum_{j\in\mathcal{N}_i}\frac{1}{v_j}} = \frac{\deg(i)}{1 + C_k v_k} \tag{4.1a}$$

$$C_k = \sum_{j\in\mathcal{N}_i \neq k} \frac{1}{v_j} \tag{4.1b}$$

The perceived intelligence factor f_{ik} represents node i's perception of node k's intelligence relative to the other members of $\mathcal{N}_i$. This information is then incorporated into the network dynamics by scaling the benefit of collaboration.[26] In this way, we now have the following network dynamics:

$$\dot{a}_{ik} = f_{ik} B'(\sigma_{ik}) - C'(d_i^{out}) \tag{4.2}$$

The factor f_{ik} helps node i indirectly incorporate information regarding its neighbors' neighbors, since v_k depends on the states of nodes in $\mathcal{N}_k$. This information is obtained at little cost; the nodes are already passing the values of their states to their neighbors during collaboration, so passing the variance v_k requires $O(\deg(k)) \leq O(N)$ extra cost of resources for node k. In sociological networks, some of this cost may be taken up by node i (the receiver of the information v_k), who may perceive the intelligence of node k through nonverbal or other indirect cues. Regardless of who takes up the cost, computing f_{ik} does not add a large degree of latency or extra burden to a network.

Properties The perceived intelligence factor exhibits desirable properties for realistically modeling intelligence of a node. First of all, the total intelligence for a node's neighbors is a conserved quantity:

$$\sum_{j\in\mathcal{N}_i} f_{ij} = \deg(i)$$

This means that the perceived intelligence of node i's neighbors is redistributed over time. Because $f_{ik} > 0$, we now know that f_{ik} does not diverge, even as $v_k \to 0$. For a finite C_k, the variation with v_k looks similar to $1/(1+v_k)$, with a maximum of $\deg(i)$ at $v_k = 0$ and an asymptote of 0 as $v_k \to \infty$. Behavior as $C_k \to \infty$ is also well defined so long as $v_k > 0$. When both $C_k \to \infty$ and $v_k \to 0$, behavior depends on the directions at which the state vectors approach the limiting values. To prevent singularities in the simulation code, we add the value 10^{-16} to each v_k before computing f_{ik}.

[26] It is reasonable to prescribe that the cost $C(d_i^{out})$ is independent of f_{ik}. Costs for node i should depend solely on the activity of node i and not on other nodes.

The perceived intelligence factor demonstrates desirable behavior for other cases of intelligence as well. Consider the case where all $v_j = a$ for $j \neq k$. Then f_{ik} becomes to the following:

$$f_{ik} = \frac{\text{deg(i)}}{1 + v_k \frac{\deg(i)-1}{a}}$$

When $v_k = a$, $f_{ik} = 1$. Therefore, the network dynamics reduce to the original network dynamics in very specialized network states. A few examples are a BCC complete graph of $N = 2$, a cycle with even N in which the states of the nodes alternate between two values, and any connected graph that has reached consensus. We require this behavior in these specialized situations, as a graph in which everyone looks equally smart should reduce to a model that does not account for intelligence. Also, $\partial f_{ik}/\partial v_k < 0$ and $\partial f_{ik}/\partial a > 0$: node k looks less smart as it drifts away from consensus with members of $\mathcal{N}_k$, and node k looks smarter as the other members of $\mathcal{N}_i$ drift away from consensus with their neighbors. Finally, the variation with degree looks like the following:

$$\frac{\partial f_{ik}}{\partial \deg(i)} = \frac{a(a - v_k)}{\left[a + v_k(\deg(i) - 1)\right]^2}$$

When node k has average intelligence, increasing the degree of i makes no difference to node k's perceived intelligence. When $v_k < a$ ($v_k > a$), increasing the number of i's neighbors makes node k look smarter (dumber). This is quite realistic, as a student who performs better than average in a class of 100 should be perceived as relatively smarter than one who performs better than average (by the same margin) in a class of 10, assuming both classes are made up of similar students.

The effect of perceived intelligence on (σ, d) dynamics is also fairly simple and intuitive: it scales the nullcline for the pair (σ_{ik}, d_i^{out}) by the factor f_{ik}, which can be seen by setting Equation 4.2 to 0 and comparing to the original formula for the nullcline (Equation 3.1). Defining the nullcline for the pair (σ_{ik}, d_i^{out}) as d_{0f} (or e_{0f} in dimensionless form), we have:

$$d_{0f}(\sigma_{ik}) = f_{ik} d_0(\sigma_{ik}) \ , \ e_{0f}(s_{ik}) = f_{ik} e_0(s_{ik}) \tag{4.3}$$

Now, different pairs of (σ, d) have different nullclines that are time-dependent, which introduces heterogeneity in the dynamics. We should expect this fact to reduce the ability of the graph to break. As seen in the reduced dynamics of Section 3.3.2, graph breakage occurs in special configurations that cross the boundary for $\dot{\sigma} = 0$ and do not return. We now revisit this system with the updated utility model.

4.1.1 Reduced System Dynamics

There are many possibilities to consider when adding initial states to the reduced system. Keeping with the spirit of the reduced system, we consider the case where one individual is smarter than all the others who have the same average intelligence. For odd N, this is easy to construct. An individual node is smartest when its state is the average of its neighbors' states:

$$\arg\min_{x_i \in \mathbb{R}} v_i = \frac{1}{\deg(i)} \sum_{j \in \mathcal{N}_i} x_j \tag{4.4}$$

Defining node i as the smart individual, we simply set x_j $(j \neq i)$ to alternate between two values for all remaining individuals in the complete BCC. This results in the desired distribution of one smart individual amongst equally average-intelligence individuals. Setting $x_j(t_0) = \pm c$ and $x_i(t_0) = 0$, we compute the following values for perceived intelligence (where i is the smart individual and nodes j and k denote other individuals):

$$f_{ji}(t_0) = \frac{(N-1)(2N-1)}{N^2 - N + 1} \tag{4.5a}$$

$$f_{jk}(t_0) = \frac{(N-1)^2}{N^2 - N + 1} = f_{ji}(t_0)\frac{N-1}{2N-1} \tag{4.5b}$$

$$f_{ij}(t_0) = 1 \tag{4.5c}$$

Note that these values do not depend on $c = |x_j(t_0)|$ due to the symmetry of the intelligence distribution. Strictly speaking, we have two choices for the smartest individual to keep the system reduced: node 1 or node 2. However, even if we allow one of the nodes in group 3 to be the smartest individual, the result is always the same steady state regardless of the direction of the small perturbation δa_{12}.

The steady-state structure is always a symmetric BCC star with the smartest individual as the leader. This result is a direct consequence of the initial condition. At t_0, the system is not close to a steady-state. The values for $\dot{\sigma}$ now appear as the following:

$$\dot{\sigma}_{mn} = f_{mn}B'(\sigma_{mn}) + f_{nm}B'(\sigma_{mn}) - 2\mu(d_m^{out} + d_n^{out}) = 4\mu\left(d_0\frac{f_{mn} + f_{nm}}{2} - \frac{d_m^{out} + d_n^{out}}{2}\right)$$

$$f_{ij}(t_0) + f_{ji}(t_0) = \frac{3N^2 - 4N + 2}{N^2 - N + 1} > 2 \text{ for } N > 3$$

$$f_{jk}(t_0) + f_{kj}(t_0) = \frac{2(N-1)^2}{N^2 - N + 1} < 2 \text{ for } N > 0$$

It is now harder (easier) for $\dot{\sigma}_{ij}$ $(\dot{\sigma}_{jk})$ to be negative in terms of average costs. Indeed, the

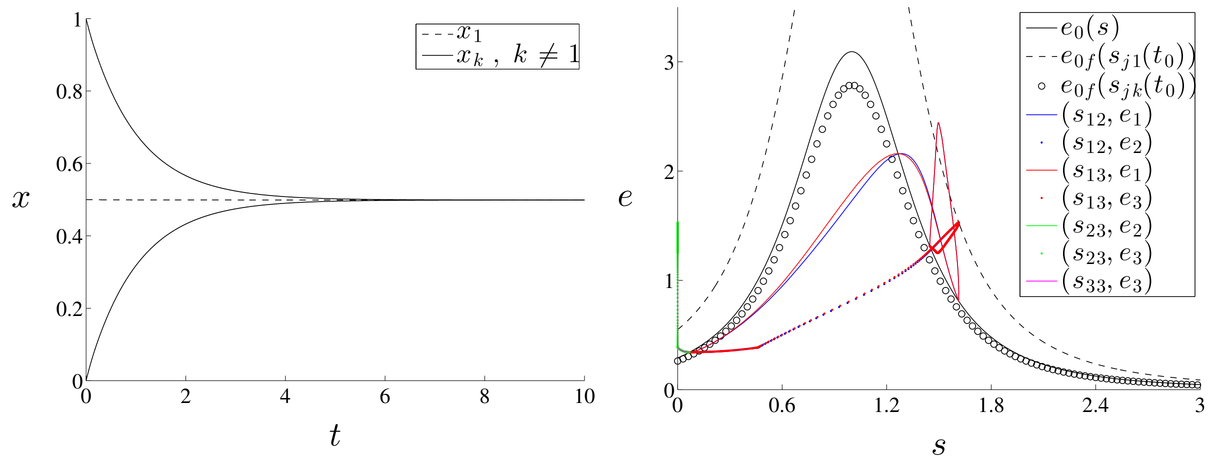

(a) State dynamics (b) Dimensionless (σ, d) dynamics

Figure 4.1. Nonlinear reduced dynamics with intelligence for $\rho = 0.65$, $\tau = 0.11$, $\mu = 1.5$, $N = 11 > N_u = 8$, and $\delta b_{12} = 0.001$. The initial state has $x_1(t_0) = 0.5$, and the other $x_j(t_))$ alternate between 0 and 1. The state dynamics are modeled without any noise to achieve a steady state in the network structure. For (b) the nullclines with intelligence are denoted by the convention in Equation 4.3 for the initial condition (i.e. $t = t_0$). The nullclines shift to $e_0(s)$ over time as the state reaches consensus. Note that node 1 remains the smartest individual for all time by the symmetry seen in (a). This is not always true for general graphs and initial conditions.

initial condition has $\dot{\sigma}_{ij} > 0$ and all other $\dot{\sigma} < 0$, which leads to the star configuration. The linearization has changed, but the dominant behavior at early times is still a reallocation of resources; the previous argument that the early behavior largely determines the steady-state is still applicable. Essentially, intelligence has separated the nullcline for the smartest individual from the others. Figure 4.1 shows the evolution with node 1 as the smartest individual and $\delta b_{12} = 0.001$. Node 1 remains the smartest individual throughout the entire time period, as can be seen by Figure 4.1(a). Indeed, $x_1 \approx 0.5$ for all time which is extremely close to the mean of the other x_j at every time step. The highly nonlinear motion for $s > 1.2$ occurs at very long times (i.e $t/\rho^2 > 20$) when the state transitions from having one leader to consensus, where all nodes have equal intelligence. Once this occurs, the dynamics reduce to the original model without intelligence, and the states fall onto the original nullcline $e_0(s)$.

We have found that perceived intelligence eliminates the breakage problem for the reduced system by introducing heterogeneity in utility functions. This makes it less probable for a (σ, d) state to cross a boundary line (such as when $\dot{\sigma} = 0$), since the nullclines no longer coincide as often during the system's evolution. In general, perceived intelligence does not prevent graph breakage, but it makes it harder to attain due to the coupling with state dynamics. We now must have a very specialized network structure *and* a very specialized state structure to have the nullclines coincide in a manner conducive to graph breakage. In other words, we have added $O(N)$ complexity to the network dynamics.

We now focus on the effects of perceived intelligence on other aspects of consensus. Our main concerns are with robustness of consensus and convergence speed, but we also examine other common time domain specifications such as overshoot for the state dynamics. All of these quantities depend on the ability to choose leaders. The perceived intelligence factor has a built-in ability to dynamically choose leaders throughout consensus, which has ramifications on consensus dynamics.

4.1.2 Dynamic Leader Selection and Overshoot

In the reduced system above, the initial leader (the smartest individual) always remained the leader. However, this is not always the case for general graphs and initial conditions. Rather, perceived intelligence implicitly encodes the ability to dynamically select leaders based on the network configuration and state. Perceived intelligence is basically a form of positive feedback: smart individuals attract collaborations and become smarter over time as their neighbors begin to follow their states through Equation 2.3a. When the graph is highly heterogeneous, either in the states x_i or the network structure (or both), neighbors who begin following a smart individual may eventually overtake this leader in intelligence and become new leaders. Additionally, the leader at any given time tracks the state of the next-smartest individuals. This drift can also result in the replacement of the leader by another node or the formation of "oligopolies" in which multiple nodes form a group in a collective leadership role. We have found all of these behaviors in simulations.

Rather than characterize the many types of leadership that are possible, we focus on the effects of dynamic leader selection on consensus characteristics. We continue to analyze the complete graph, as this allows us to plot the N values of absolute intelligence $1/v_k$ as a proxy for plotting all (typically $O(N^2)$) f_{ij} values over time. We keep the initial condition for $\mathbf{x}$ the same as with the dynamics in the previous section, with node 1 as the clear leader and all other nodes as average in intelligence. This construction allows for comparisons with previous results. Finally, we add Gaussian noise to the initial conditions for the network structure. For illustrative purposes, we use the same system of $N = 11$ but change the initial condition so that $a_{ij}(t_0) = \mathcal{N}(10, 3^2)$, which represents a gross initial over-investment of resources.

State Dynamics Figure 4.2 shows the state dynamics for the resulting network both with and without intelligence. We also compare both of these conditions with a static network

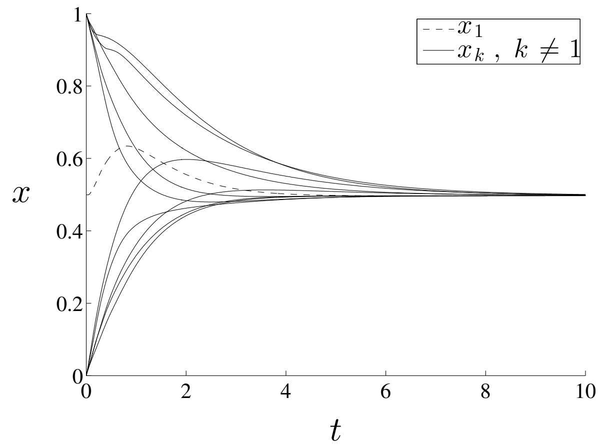

(a) State dynamics without intelligence

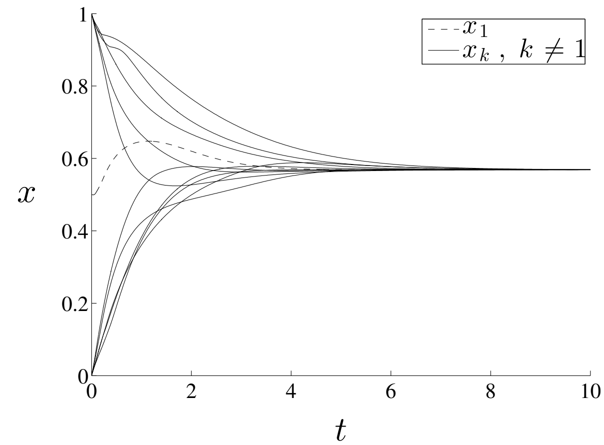

(b) State dynamics with intelligence

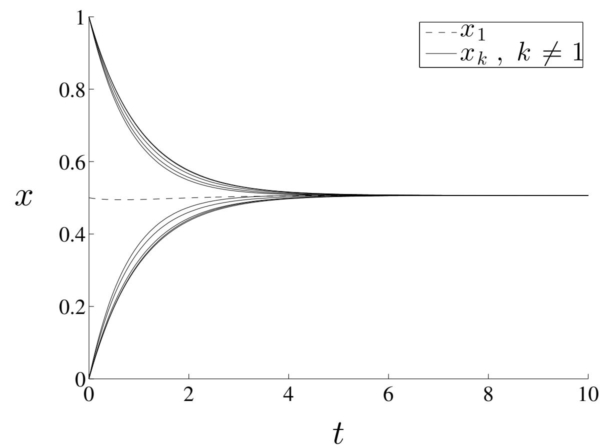

(c) State dynamics with static network

Figure 4.2. State dynamics with noisy initial network conditions: $\rho = 0.65$, $\tau = 0.11$, $\mu = 1.5$, $N = 11 > N_u = 8$, and $a_{ij}(t_0) \in \mathcal{N}(10, 3^2)$. The state dynamics are simulated without noise and initial conditions set $x_1 = 0.5$ and the other x_j alternate between 0 and 1. For (c) the network is static and remains at its initial condition for all time. This configuration results in speedy convergence, but it comes at extremely high costs for all individuals since they over-invest into collaborations. We account for such costs in Section 4.3.

stuck at the initial condition. Intelligence has a smoothing effect for the state dynamics. We can define the overshoot (OS) of a particular state x_i as the following:

$$OS = \frac{|x_i(t_1) - x_i(t_0)| - |x_i(t_f) - x_i(t_0)|}{|x_i(t_f) - x_i(t_0)|} \tag{4.6}$$

where $x_i(t_f)$ is the final value at consensus and $x_i(t_1)$ is the value of the state at its maximum distance from the initial condition.[27] Even though the overshoot for node 1 is the greatest, we

[27] The absolute values account for the fact that the consensus state may be on the opposite side of the initial condition as the direction in which the maximum deflection occurs.

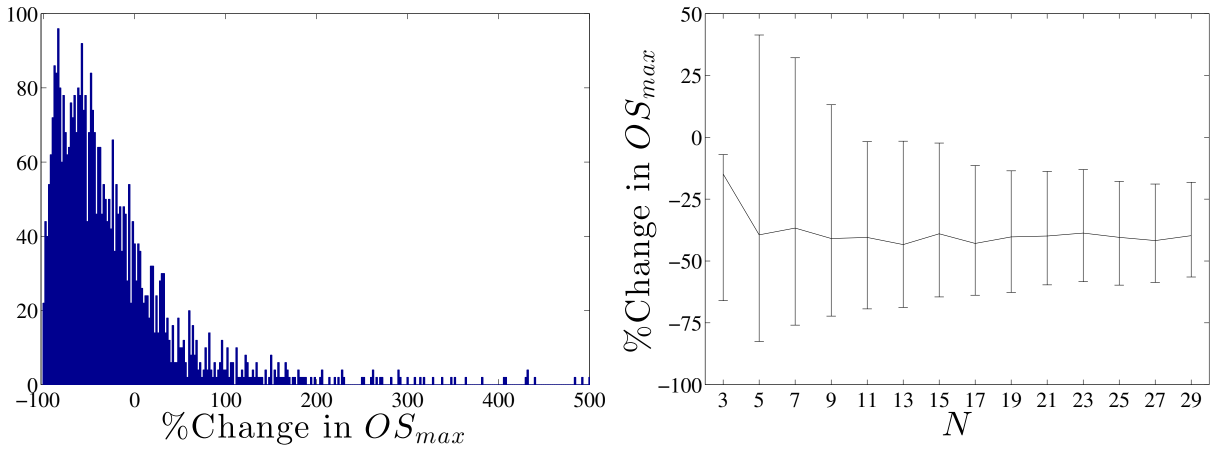

(a) Change in overshoot, $N = 11$

(b) Change in overshoot with N

Figure 4.3. (a) Histogram of percent change in overshoot with intelligence (shown with a bin width of 2). 4000 simulations with the conditions of Figure 4.2 were conducted, and the maximum overshoot value for one of the average-intelligence nodes was compared between the cases with and without intelligence. A value between -100 and 0 indicates a reduction in maximum overshoot for the state dynamics (see Equation 4.7). The median is -43 and the distribution is highly skewed towards overshoot reduction.
(b) % Change in OS_{max} vs. N. At each value of $N \in \{3, 5, 7, ..., 29\}$, 1000 simulations were carried out with the same Gaussian initial conditions as in Figure 4.2. The centerline is the median, and the error bars span the first to third quartiles of the data. The distributions are highly skewed, so these measures are more suitable than the mean and standard deviation.

neglect node 1 because its overshoot is unreasonably magnified by the proximity of consensus to its initial state. Considering only the average-intelligence nodes, overshoot drops from 19.4% to 10.3% with intelligence for this case, a roughly 47% reduction. Figure 4.3(a) shows similar results for this simulation repeated 4000 times with initial conditions chosen from the same Gaussian distribution. The change in overshoot was calculated according to the following formula:

$$\% \text{ Change in } OS_{max} = 100\frac{OS_{max,intel} - OS_{max}}{OS_{max}} \tag{4.7}$$

where OS_{max} is the maximum overshoot for one of the average intelligence nodes without intelligence dynamics, and $OS_{max,intel}$ is the maximum overshoot for one of the average intelligence nodes with intelligence dynamics activated. The resulting distribution is highly skewed towards reducing maximum overshoot, with a median value of -43%. Figure 4.3(b) repeats these statistics with $N \in \{3, 5, 7, ..., 29\}$. The resulting distribution (plotted with the median spanned by the first and third quartiles) decreases in spread with N and approaches an asymptote of roughly -40 ± 10, indicating a clear reduction in maximum overshoot with intelligence but definitely not an elimination of overshoot altogether.

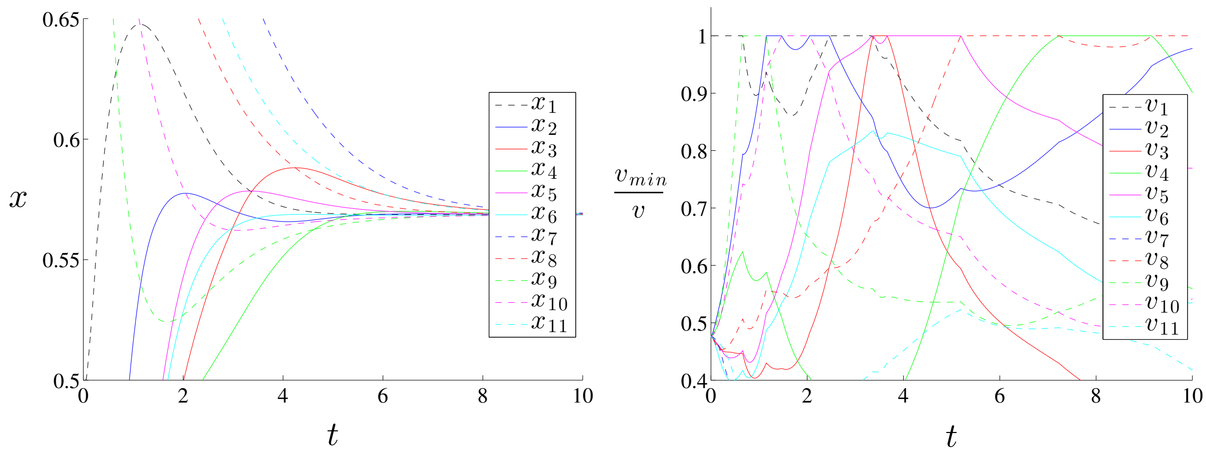

(a) State dynamics (b) Normalized absolute intelligence dynamics

Figure 4.4. Dynamic leader selection and overshoot. We show the same simulation from Figure 4.2, tracking both the state dynamics (with intelligence) and the absolute intelligence. The leader always has $v_{min}/v = 1$. Exchange of leadership occurs when the current leader's state crosses with another node's state. This can only occur if the leader drifts away from consensus, which is the same condition for overshoot of the leader's state. At long times, ($t > 25$ in this case) the states converge to all have the same absolute intelligence (which is not shown here).

The intuition behind this smoothing effect is the following: with intelligence, the dynamics adapt accordingly to new leadership, thereby preventing the nodes from following individuals who are no longer close to consensus. Without intelligence, the network evolution does not care about the state dynamics, so the probability that the network evolution and state dynamics converge in a perfectly complementary manner is very small. Whenever the network moves in a direction that does not perfectly coincide with the motion of the consensus state, overshoot results. In this way, intelligence reduces the amount of overshoot compared to the original utility model. Compared to the static network, however, intelligence still produces a significant amount of overshoot in the state dynamics. Additionally, intelligence induces a different type of overshoot, one that is associated the graph structure rather than the state dynamics. This phenomenon is directly related to dynamic leader selection.

To see the relationship between network overshoot and dynamic leader selection, we need to track intelligence. Figure 4.4 shows the normalized absolute intelligence $(1/v_k(t))$ for each node over time with the same initial condition as in Figure 4.2. The normalization is chosen such that the smartest individual always sits at 1, thereby making the normalization factor $\max_k 1/v_k(t)$ which corresponds to $\min_k v_k(t) := v_{min}(t)$.[28] We also show the state dynamics as in Figure 4.2 with every node labeled individually. Leadership exchange occurs when the current leader's state drifts far enough that its state meets and then exceeds the state of

[28] Note that this value changes with time.

another individual. This can be seen by the numerous exchanges of leadership from node 1 to 9, 9 to 2, 2 to 10, etc. The major point of concern regarding network overshoot is the dramatic refocus of leadership even as the states reaches consensus. Although the frequency of leadership exchange slows down as the system asymptotically approaches consensus, the former leaders often exhibit enormous drops in intelligence (such as node 1), and other nodes exhibit similarly dramatic rises in intelligence (such as node 4). These swings in intelligence are undesirable because they cause the graph structure to make dramatic changes even as the states reach consensus: a node that is at one point a focus of attention quickly becomes a fringe follower as its intelligence reduces. This behavior is related to state overshoot because the nodes that start following a certain leader swing too quickly towards its state, thereby overshooting it. As with any system, overshoot results from insufficient damping. In this case, we are not damping the response to perceived intelligence.

Robustness and Convergence Speed The dramatic changes in graph structure due to the undamped responses to intelligence have direct consequences on the robustness and speed of convergence for the network. Figure 4.5 compares the speed of convergence and robustness of both network models. For both metrics, the case with intelligence has better average values,[29] but intelligence also creates wildly oscillatory behavior. This can be attributed to the swings in network structure due to hasty leader selection and overshoot. In particular, note that the local minima (maxima) of λ_{min} correlate with the local maxima (minima) of H_2. Recall that the homogeneous complete graph has the best convergence speed and smallest H_2-norm. Then, loosely speaking, the rising (falling) of λ_{min} and falling (rising) of H_2 coincide with the state moving towards being more (less) complete. In this way, a local extremum on either graph corresponds to a swing in behavior for network evolution. Sharp points correspond to edges vanishing or reappearing. Comparing these plots with Figure 4.4(b), we see that the swing in behavior follows an exchange in leadership by a time delay of between 0.25 and 0.75 units of time, which is $O(\rho^2)$ in this case. This value makes intuitive sense because the network evolves on a timescale of ρ^2, so it should take $O(\rho^2)$ time to respond to leadership exchange. The local maxima of λ_{min} occur at the half-way transition points between following one leader to following the new leader, the point at which the graph "fills in" its edges and (locally) becomes as complete as possible before swaying towards the new leader.

To see these oscillations more clearly, consider the simple case of two symmetric BCC stars with $N = 3$, one with node 1 as the leader, and another with node 2 as the leader. In Figure 4.6, we show the weighted average of these two graphs, with a weight of α for the graph

[29] We aim to maximize convergence speed and minimize the H_2-norm.

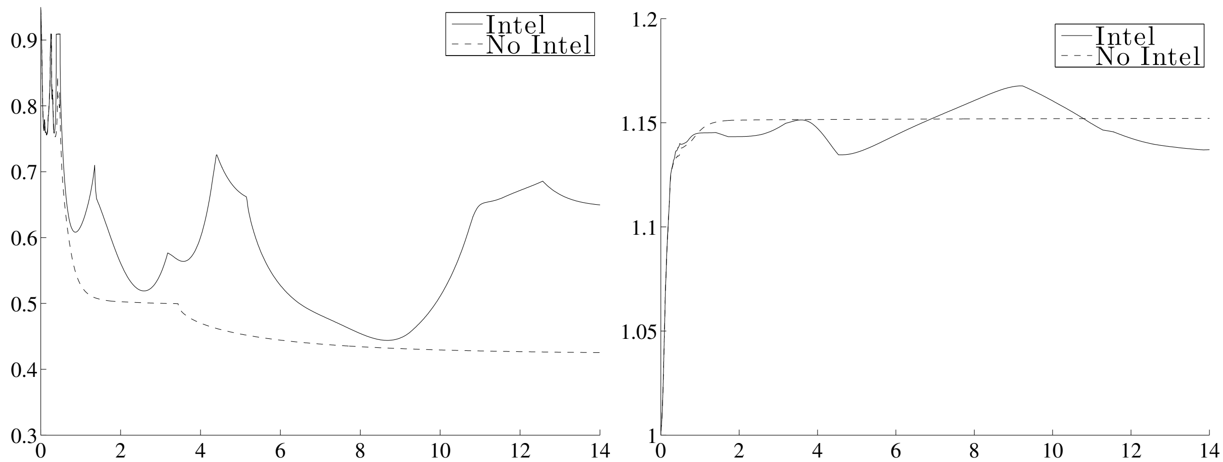

(a) $\lambda_{min}/\lambda^*_{min}$ vs. t (b) H_2/H_2^* vs. t

Figure 4.5. Robustness and convergence speed with and without intelligence for the same case as that in Figure 4.2. With intelligence, the dramatic swings in graph structure due to hasty leader selection and overshoot result in oscillations for both metrics.

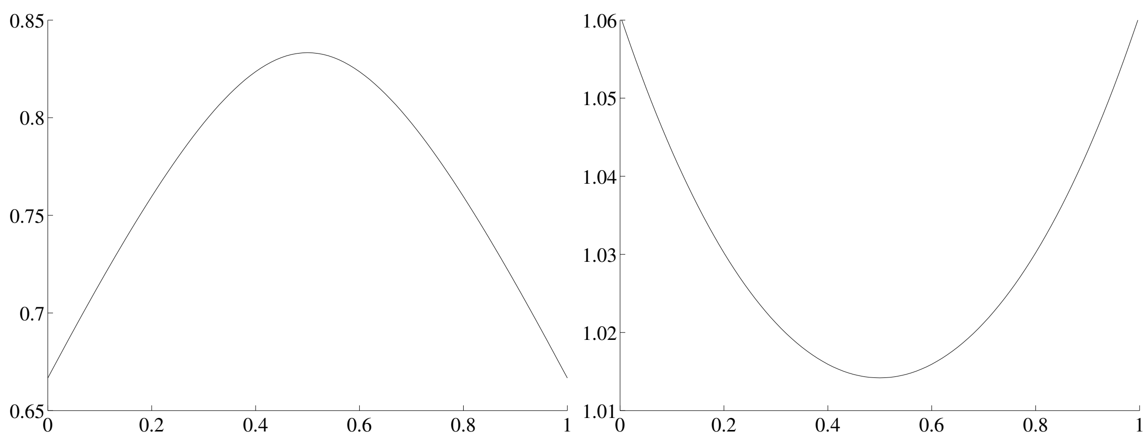

(a) $\lambda_{min}/\lambda^*_{min}$ vs. α (b) H_2/H_2^* vs. α

Figure 4.6. Convergence speed and robustness for the weighted average of two symmetric BCC stars ($N = 3$) that are at equilibrium. The coefficient α is the weight of the graph with node 1 as the leader, and $1 - \alpha$ is the weight for the graph with node 2 as the leader. The extremum corresponds to the point at which the graph becomes as complete as possible, the halfway point between both stars. Changing α mimics the effect of leadership exchange and the resulting swings in graph structure for a network with perceived intelligence dynamics.

with node 1 as the leader, and $1 - \alpha$ as the weight of the other. The maximum λ_{min} and minimum H_2 occur when the graph is as complete as possible, the midway point between the two stars. This figure mimics the effect of leadership exchange in a simplified way. The real system does not necessarily swing all the way to a star before exchanging leaders again. In addition, the smooth behavior is often interrupted by an edge vanishing or reappearing, which creates the kinks in Figure 4.5(a).

Main Deficiency of Perceived Intelligence The overall drawback to the perceived intelligence model is that it lacks a damping factor. This results in an overdramatic response to changes in state, as illustrated by overshoot and dynamic leader selection. The ability to dynamically update leaders is not intrinsically bad. Indeed, it is very useful in cases where we may have node failures, such as when a sensor fails on a robot amidst a mobile sensing network. In our present case, however, the system hastily forgets about nodes who were initially intelligent, and new leaders arise simply because of overshoot. This "mob mentality" can result in the selection of individuals who are not necessarily the best choices as leaders. Such suboptimal performance can be amended if we allow agents to not only track the current intelligence of their neighbors but also retain a history of perceived intelligence.

4.2 Perceived Intelligence with History

In the model system of students collaborating on a homework assignment, we would not expect students to merely track current intelligence when considering collaborations. Rather, we expect a student to take into account the past history of a neighbor's actions, as dictated by his/her perceived intelligence over time. This intuition motivates the model of an estimation scheme in which nodes estimate the perceived intelligence of their neighbors based on a certain "sliding window" of past states.

Literally storing the previous intelligence states of neighbors is both inefficient and unrealistic. In the spirit of recursive estimators such as the Kalman filter, we define a recursive discrete-time update rule for perceived intelligence. In this new model, we define the variable $\hat{f}$ as the estimated value for perceived intelligence, as opposed to f which is the actual value. Then the update rule is written as follows in discrete time steps n:

$$\hat{f}_n = \beta \hat{f}_{n-1} + (1-\beta) f \ , \ \beta \in [0,1] \tag{4.8}$$

where β defines the weight with which we take the past estimation of intelligence. As $\beta \to 1$, our window extends to the beginning of time, whereas $\beta \to 0$ reduces to the model without history. To convert this model to continuous time, we first rewrite Equation 4.8 as the following:

$$\frac{\hat{f}_n - \hat{f}_{n-1}}{\Delta t} = \frac{1-\beta}{\Delta t}\left(f - \hat{f}_{n-1}\right) \tag{4.9}$$

where Δt is the time step between the discrete states. Note that $\beta = \beta(\Delta t)$. Indeed, as Δt gets large, we prescribe $\beta \to 0$: since the state could have changed dramatically between time steps, we are better off just following the new value for intelligence. Similarly, as $\Delta t \to 0$ we prescribe $\beta \to O(1-\Delta t)$: the non-homogeneous term adds very little to our current estimate, so it does not need a large weight. This reasoning justifies the fact that $\frac{1-\beta}{\Delta t}$ remains finite as $\Delta t \to 0$. Then we can write the continuous time dynamics as:

$$\dot{\hat{f}} = \gamma\left(f - \hat{f}\right)\ ,\ \gamma := \frac{1-\beta}{\Delta t} \in [0, \infty) \tag{4.10}$$

As $\gamma \to 0$, the nodes maintain a history over all time and do not look at the new values for intelligence (i.e. the estimates remain fixed at the initial values). As $\gamma \to \infty$, the nodes do not maintain history at all. Therefore, γ governs the amount of damping for the response of each node to perceived intelligence. The updated network dynamics now look like:

$$\dot{a}_{ik} = \hat{f}_{ik} B'(\sigma_{ik}) - C'(d_i^{out}) \tag{4.11a}$$

$$\dot{\hat{f}}_{ik} = \gamma\left(f_{ik} - \hat{f}_{ik}\right) \tag{4.11b}$$

where $\hat{f}_{ik}(t_0) = f_{ik}(t_0)$. Thus, the overall system now contains state dynamics ($\dot{\mathbf{x}}$), network dynamics ($\dot{a}_{ik}$), and intelligence dynamics ($\dot{\hat{f}}_{ik}$).

4.2.1 Effects of History on Network Dynamics

We have introduced history as a method to dampen the swings in graph structure. As we illustrated above, these swings can be tracked by oscillations in the speed of convergence and the robustness over time. Figure 4.7 shows the effect of varying γ on simulations using the same initial conditions as those in Figure 4.2. As predicted, large histories ($\gamma < O(1)$) reduce both the amplitude and frequency of oscillations. For short histories ($\gamma > O(1)$), the variations approximate those in Figure 4.5 without history, with a short time delay that diminishes as γ increases. At $\gamma = O(1)$, the oscillations are delayed by $\Delta t = O(1)$ but occur at roughly the same frequency, and the amplitude of oscillations is roughly the same if not slightly larger at some points. These results are simple to explain by noting how the history dynamics couple with the intelligence and network dynamics. According to Equation 4.11b, the estimated intelligence dynamics evolve on the timescale $1/\gamma$. For small γ, the variations in leadership and intelligence are extremely damped, whereas for large γ they approximate the undamped case without history. At $\gamma = O(1)$, the estimated intelligence resonates with changes in leadership. We know by Figure 4.4 that the changes in leadership occur on a timescale that is $O(1)$ for intelligence without history, so this must also be the case for history with $\gamma = O(1)$. Indeed, for $\gamma = O(1)$, $\hat{f}$ follows f with an $O(1)$ time delay (by Equation 4.11b), which is short enough to respond to all changes in leadership. At the same time, however, this delay is long enough that the graph can make slightly more dramatic swings: since $\hat{f}(t_0) = f(t_0)$ and the network takes an extra $O(\rho^2)$ delay to respond to changes

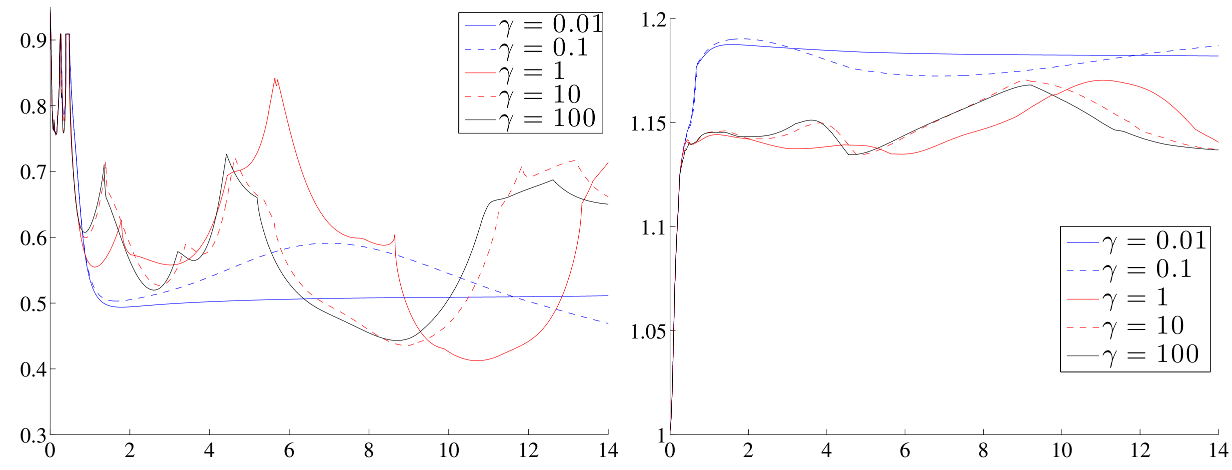

(a) $\lambda_{min}/\lambda^*_{min}$ vs. t (b) H_2/H_2^* vs. t

Figure 4.7. Robustness and convergence speed with history for the same case as that in Figure 4.2. For $\gamma < O(1)$, both the amplitude and frequency of oscillations is diminished. For $\gamma > O(1)$, the metrics approximate those without history with a time delay that diminishes as γ grows. At $\gamma = O(1)$, history dynamics resonate with state dynamics such that the frequency and magnitudes are roughly the same (and sometimes slightly larger), with an $O(1)$ time delay.

in $\hat{f}$, we induce the capability to swing slightly more in a certain direction for $O(\rho^2)$ time. Based on this analysis, it is evident that we should have $\gamma < O(1)$ for history to provide significant benefits in damping network swings.

In addition to damping the response to leadership exchange, we expect history to also slow down the rate of leadership exchange. Since $\gamma > O(1)$ essentially mimics intelligence without history, we consider only $\gamma \leq O(1)$. Figure 4.8 compares the absolute intelligence dynamics for the extremes $\gamma = 0.01$ and $\gamma = 1$. For $t > 2$, the former case has much less frequent leader exchange, whereas the latter case looks extremely similar to the situation without history in Figure 4.4(b). This is expected by the analysis of $\gamma = O(1)$ above. However, the frequency of leader exchange for $\gamma = 0.01$ at $t < 2$ is troubling, since we we clearly want node 1 to remain the leader for all time. This problem is related to the model's inability to damp the *leader's* overshoot, which we discuss next.

4.2.2 Effects of History on State Dynamics

As opposed to its ability to easily damp swings in graph structure, history is not able to uniformly damp overshoot in state dynamics. Figure 4.9(a) shows the same simulation conditions of Figure 4.2 with $\gamma = 0.01$. The extremely large overshoot for one of the average intelligence nodes occurs due to the motion of node 1. With $\gamma = 0.01$, perceived intelligence estimates are essentially fixed and the values given by Equation 4.5 hold for $O(1/\gamma)$ time.

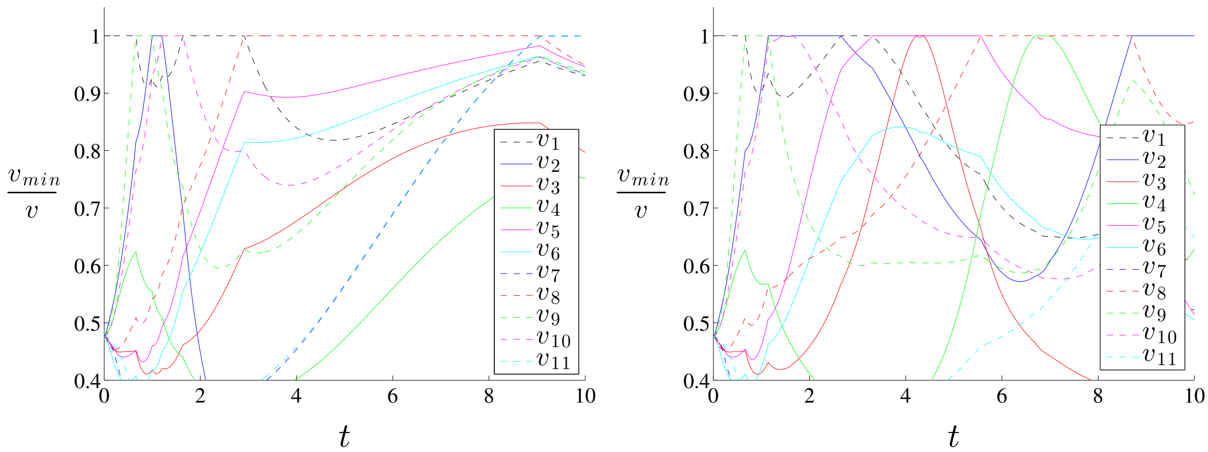

(a) $\gamma = 0.01$ (b) $\gamma = 1$

Figure 4.8. Absolute intelligence dynamics with history. Leadership exchange is less frequent with longer histories.

Thus, node 1 is considered to be the leader for a long period of time. Additionally, node 1's beliefs about all neighbor's are also fixed; its model is reduced to the one without intelligence since it believes all neighbors have equal intelligence. Therefore, node 1 moves according to the original utility model without intelligence, and other nodes follow its motion very closely, resulting in large overshoot for those nodes j whose value for $a_{j1}(t_0)$ happened to be much larger than all other $a_{jk}(t_0)$ $(k \neq 1)$. The magnitude of overshoot decreases with increasing γ, as shown in Figure 4.9(b). Here, we track the change in maximum overshoot with intelligence by the following formula:

$$\% \text{ Change in } OS_{max,intel} = 100 \frac{OS_{max,hist} - OS_{max,intel}}{OS_{max,intel}} \tag{4.12}$$

where $OS_{max,hist}$ is the overshoot with the given value of γ for history and $OS_{max,intel}$ still represents the overshoot for intelligence without history. Despite the large spread at small γ, the lower quartile only falls into a mode of reduction at $\gamma = 1$, and we require $\gamma < O(1)$ for reasonable damping of network swings.

We have found that the model of perceived intelligence with history fails to properly account for the motion of a true leader. What we require in this case is for the leader to remain in a leadership position. This can be achieved in multiple ways, such as having the leader smoothly equalize its investment across all neighbors or simply remain fixed at its current state. However, it is important to emphasize that the case of a single global leader is very specialized. A realistic network such as a small-world graph is not complete, and multiple nodes act as local leaders; having a single global leader is improbable over such a graph.

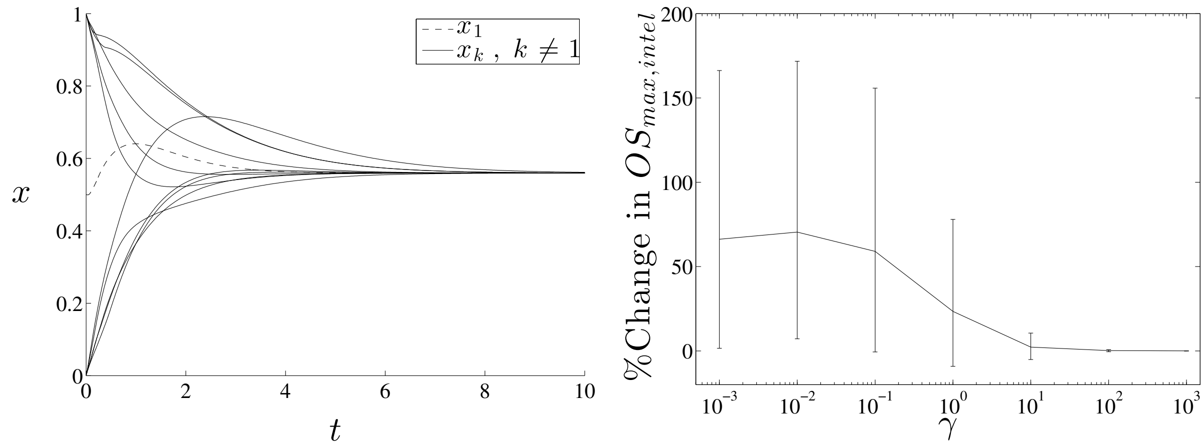

(a) State dynamics with $N = 11$, $\gamma = 0.01$ (b) Change in overshoot with intelligence over γ

Figure 4.9. (a) State dynamics for the exact initial condition as in Figure 4.2 and $\gamma = 0.01$. Since the intelligence estimates are essentially fixed, the motion of node 1 causes overshoot for any average-intelligence nodes who happen to follow it strongly in the initial condition, leading to dramatic overshoot.
(b) Change in overshoot with history. 4000 simulations were conducted at each γ with $N = 11$ and Gaussian initial conditions as in Figure 4.2. The median is shown spanned by the first and third quartiles in error bars. The movement of node 1 causes worse overshoot with longer histories: other nodes believe it is still smart despite leadership exchange taking place.

The overall problem we have posed is designed to work in these situations, where a global leader is not immediately discernible. In the complete graph above, we saw that any motion of the global leader was sub-optimal for reaching consensus. However, the motion of a local leader towards its neighbors may improve performance depending on the graph and state structures. In this light, we continue to evaluate the utility model with perceived intelligence and history as a distributed heuristic for reaching consensus. This analysis will also yield further methods by which to address problems with the dynamics of (global or local) leaders.

4.3 Evaluations of Network Performance

Perceived intelligence and history sophisticate the original utility model by attempting to overcome problems such as network breakage and dramatic swings in network structure. We now evaluate the performance of our model with respect to consensus dynamics. Calculating the speed of convergence and robustness is simple for a static network. Since $\hat{L}$ is fixed, the values λ_{min} and H_2 are good metrics. However, both of these values change with time in an evolving network, and it is not enough to just track the instantaneous values over time.

Effective Speed At first glance, it might appear adequate to merely keep a running average of $\lambda_{min}(t)$. However, this does not account for the fact that slower λ_{min} delays the time at which consensus is reached. Instead, we define the effective speed as:

$$\lambda_{min}^{e} = \frac{1}{t_s} \tag{4.13}$$

where t_s is the value at which the states have reached some threshold near the final consensus value. When simulating with noise, the consensus value changes with time, so it can be taken as the average of the states at long times. Also, the threshold we choose should depend on the magnitude of $\boldsymbol{\xi}(t)$, the noise in the state dynamics (Equation 2.3a). As we have been simulating without noise, we can choose an arbitrarily small threshold: we choose 99% so that t_s is the time after which all states remain within 1% of the final consensus state (which is fixed when there is no noise). The value of t_s is analogous to the settling time of a conventional SISO system.

Effective Robustness Speed is primarily about *convergence* characteristics, whereas robustness is a metric for *consensus*. In this way, looking at cumulative snapshots of $H_2(t)$ while the network is still in its transient phases does not make physical sense for evaluation. Rather, we measure the effective robustness using a snapshot value for the effective H_2-norm:

$$H_2^e = H_2(t_s) = H_2\left(\frac{1}{\lambda_{min}^{e}}\right) \tag{4.14}$$

Effective robustness (proportional to the inverse of H_2^e) characterizes the response of the system to further perturbations once it has settled to a near-consensus state.

In Figure 4.10, we calculate the effective speed and robustness over different values of γ, compared with the cases of intelligence without history and no intelligence. In these simulations, we remove the constraint of having only one leader and instead have $x_i(t_0) = \mathcal{N}(0.5, 0.2^2)$. The comparisons are as follows:

$$\%\text{Change from } *^e = 100\frac{*_{hist}^{e} - *^e}{*^e} \tag{4.15a}$$

$$\%\text{Change from } *_{intel}^{e} = 100\frac{*_{hist}^{e} - *_{intel}^{e}}{*_{intel}^{e}} \tag{4.15b}$$

where $*$ is a placeholder for either λ_{min} or H_2. Compared to the system without intelligence, the average speed increases as history decreases (γ increases) and the system approaches intelligence without history. Equivalently, intelligence without history performs better than systems with long histories. This result illustrates the notion that despite the benefit of

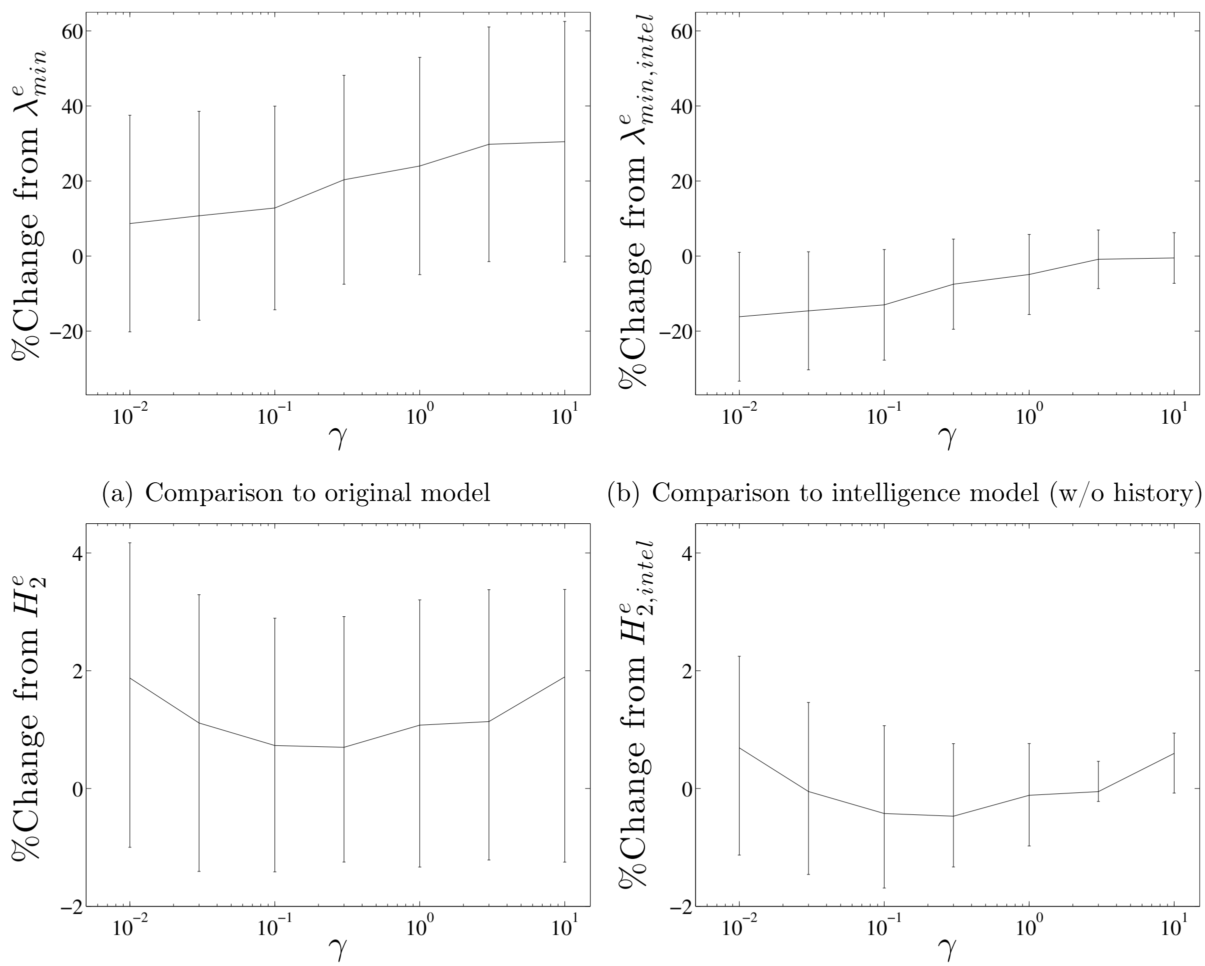

(a) Comparison to original model

(b) Comparison to intelligence model (w/o history)

(c) Comparison to original model

(d) Comparison to intelligence model (w/o history)

Figure 4.10. Effective speed and robustness. 500 simulations were carried out at each value of $\gamma \in \{0.01, 0.03, 0.1, 0.3, 1, 3, 10\}$ with $\rho = 0.65$, $\tau = 0.11$, $\mu = 1.5$, $N = 11$, $a_{ij}(t_0) \in \mathcal{N}(10, 3^2)$, and $x_i(t_0) \in \mathcal{N}(0.5, 0.2^2)$. These are compared with the cases of no intelligence and intelligence without history. For (a) and (b) we plot the mean spanned by the standard deviation because the distributions are relatively symmetric without outliers. For (c) and (d) we plot the median spanned by the first and third quartiles to resist slight skew and outliers.

reducing swings in graph structure, long histories simply prevent nodes from responding to changes in leadership. This leads to a slower effective speed than in a more adaptable network. The effective robustness shows very little change when compared to either the system without intelligence or the system with intelligence but no history. Nevertheless, we see a slight minimum at $\gamma = O(10^{-1})$, indicating that despite the lag induced by history, the resulting damping of graph swings *may* help the network attain better absolute performance to perturbations. We are careful not to read too much into this data due to the fact that 0% change is within the margin of error for both plots. However, what appear as only marginal improvements in performance are actually much better once we consider network costs.

Accounting for Network Costs Neither λ^e_{min} nor H^e_2 take into account the costs associated with consensus. If absolute speed and robustness were all that mattered, we would always prefer a static homogeneous complete BCC over any other network. Also, we would not care about the magnitudes of a_{ij} since $\hat{L}$ is normalized. This is unrealistic in two ways. First, we should penalize "completeness" in systems, since a star configuration is generally more efficient than a complete graph. Also, we should penalize the absolute magnitudes of a_{ij}, since a homogeneous complete BCC with small a_{ij} is obviously more efficient than one with large a_{ij}. We can account for both of these factors at the same time by defining the network cost factor:

$$NC(t) = \left(\frac{\sum_{i,j} a_{ij}(t)}{N^2 - N}\right) \sum_{i,j} I(a_{ij}(t)) \tag{4.16a}$$

$$I(a_{ij}) = \begin{cases} 1 & \text{if } a_{ij} > 0 \\ 0 & \text{if } a_{ij} = 0 \end{cases} \tag{4.16b}$$

The first term is the average value of a potentially nonzero entry in A, since there can be at most $N^2 - N$ nonzero links. This term penalizes the absolute magnitudes of a_{ij}. The second term is an indicator variable that counts the nonzero elements in A, which penalizes systems that are more complete than others. For a complete graph, $NC(t) = \sum_{ij} a_{ij}$, whereas for a star, $NC(t) = 2\sum_{ij} a_{ij}/N$. Now we can define the cost-normalized effective speed and H_2-norm:

$$\lambda^{cn}_{min} = \frac{\lambda^e_{min}}{NC(t_s)} \tag{4.17a}$$

$$H^{cn}_2 = NC(t_s)H^e_2 \tag{4.17b}$$

The appropriate reference values for the homogeneous complete graph at equilibrium are λ^{cn*}_{min} and H^{cn*}_2. Note that these reference values depend on the magnitudes of a_{ij} at equilibrium for the homogeneous complete graph, as opposed to λ^*_{min} and H^*_2, which only depend on N. A major benefit to absorbing $NC(t)$ into the metrics for speed and robustness is that we do not have to examine the graph structure explicitly to evaluate performance, despite the fact the network has changed over time. This allows for analysis that is normally done over static networks to easily extend to evolving networks.

Figure 4.11 shows the plots of Figure 4.10 updated to include $NC(t_s)$ ($*^{cn}$ replaces $*^e$ in Equation 4.15). The plots for cost-normalized effective speed are similar to those of Figure 4.10, but they are shifted so that the performance of the original model without intelligence is now much worse than those with intelligence. Again, we see that the graph tends to be more efficient with intelligence because nodes quickly find and follow smart leaders more intensely. The cost-normalized effective robustness plots also experience this separation; the original

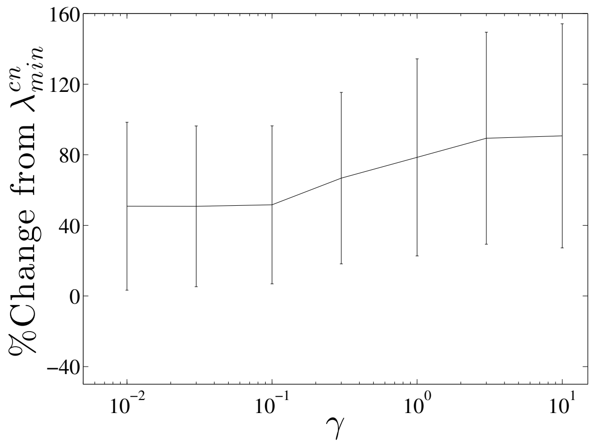

(a) Comparison to original model

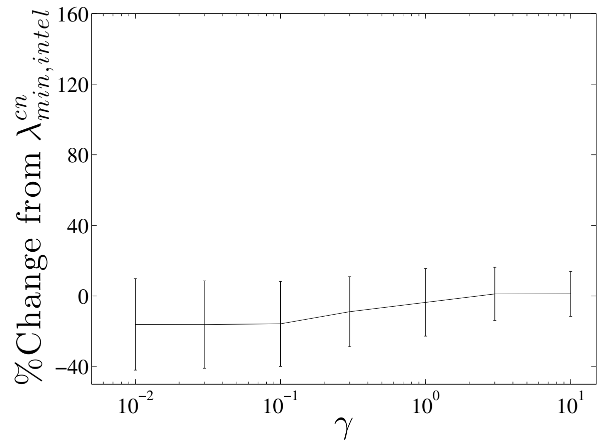

(b) Comparison to intelligence model (w/o history)

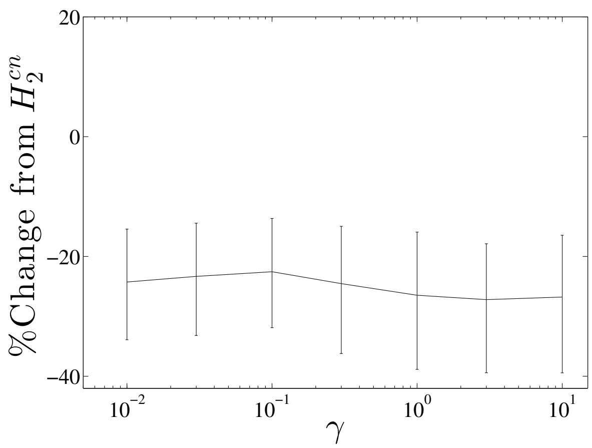

(c) Comparison to original model

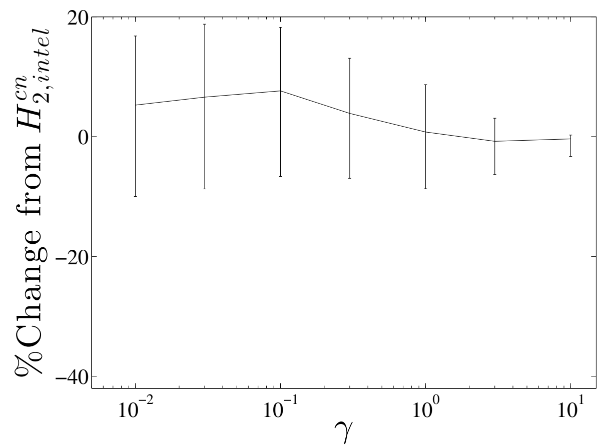

(d) Comparison to intelligence model (w/o history)

Figure 4.11. Cost-normalized effective speed and robustness. The results from Figure 4.10 are shown with the cost-normalization factors $NC(t_s)$ included. Now, the original utility model without intelligence pales in comparison to the other models, as cost-normalization has induced a separation of the models for both metrics. Intelligence without history performs slightly better on average than intelligence with history, both in speed and robustness, but these two models are roughly equivalent within the margin of error of the data spread.

model without intelligence is much worse than the models with intelligence. Although the trends for the average value are now inverted, more simulations are needed to reduce the spread and make more definitive statements. Without more data, we can only comment on the fact that including cost-normalization has clearly separated the models with intelligence from the original model in both speed and robustness. Coupling this result with the analysis of hidden modes and graph breakage in reduced systems, we have good reason to discard the original utility model and consider only the intelligence models (with or without history) for further analysis.

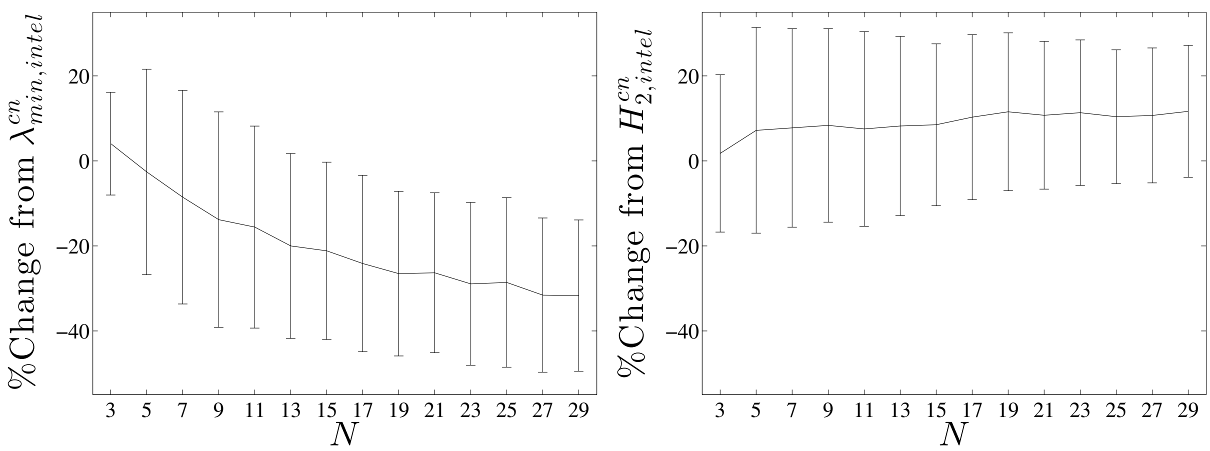

(a) Cost-normalized effective speed (b) Cost-normalized effective robustness

Figure 4.12. Change in cost-normalized effective metrics with N. The conditions of Figure 4.10 were repeated with $\gamma = 0.1$ and $N \in \{3, 5, 7, ..., 29\}$ (with 500 simulations at each N). Despite the large spread of the data, we see that performance degradation increases with N. Since the graph begins as a complete network, N approximates $\deg_t(i)\ \forall\ i$ in the transient before the system settles. This trend corroborates the intuition behind making γ depend on $\deg_t(i)$.

The Need for Degree-Dependent Histories It is clear that the models with intelligence perform better than the model without intelligence. Also, history degrades performance for systems that have a single global leader. For systems without a clear leader, current analysis cannot yet make definitive claims regarding the effects of history. At this point, history appears to change performance very little within the margins of deviation for the data presented in graphs that begin as complete with $N = 11$.[30] It is also important to consider changes in performance with changing values of N.

Going back to the situation of students solving a homework problem, we expect the level of history that each student tracks to depend on his/her centrality; it is much harder to keep track of long histories for 100 neighbors than just 5 neighbors. This intuition predicts that we should have γ increase with the instantaneous outdegree of a node, $\deg_t(i)$.[31] Far more than a prediction, this hypothesis is validated by data. Figure 4.12 shows the change in performance for networks with $\gamma = 0.1$ and $N \in \{3, 5, ..., 29\}$. Because the initial network conditions ($a_{ij}(t_0)$) represent a complete graph, N acts as a rough approximation for $\deg_t(i)\ \forall\ i$ for the transient before the system settles. This allows us to easily test our hypothesis without needing to track $\deg_t(i)$ for all individual nodes over time. Despite the large

[30] The graphs always settle into a sparse network structure as indicated by the values for λ_{min} and H_2 in Figure 4.7.

[31] This value is similar to $\deg(i)$ except that $\deg_t(i)$ changes with time and is a directed measure. Specifically, $\deg_t(i)$ counts the number of nonzero $a_{ij}(t)$.

spread for the data, the trend of increased performance degradation with N is undeniable. We take this observation as sufficient motivation for requiring that $\gamma = \gamma(\deg_t(i))$. With this sophistication, the problems we encountered with a single global leader and history will also be assuaged; leaders will keep smaller histories than other nodes, thereby reducing the overshoot problems we encountered in Section 4.2.2. Coupled with the changes to $\dot{a}_{ij}$ we proposed for leaders in Section 4.2.2, this modification will surely improve the performance of our utility model in regards to both speed and robustness.

5 Conclusions and Further Work

This study focused on developing a distributed consensus protocol for adaptive multi-agent systems. The stated requirement for such a protocol was that it should be *locally adaptive*, i.e. changes to the network structure should not require knowledge of global network properties. A utility maximization approach was adopted, whereby local updates to network structure resulted from the optimizations of individual utility functions by all nodes in the network. Beginning with a utility function inspired by economic and sociological models, various sophistications were added to this function to achieve desired behaviors in state and network dynamics. In a larger sense, we have developed a fairly robust model capable of adapting to numerous environments and applications through the specification of only a few parameters.

The protocol in its current form evolved quite organically from the model presented in [5]. Dimensional analysis of this model revealed the network timescale as ρ^2, an important characteristic which depends solely on the point of inflection for the benefit function in utility; this factor remains the network timescale for the current protocol. Further analysis of the original model revealed two major shortcomings. The first had to do with redundancy: the model of [5] fails to respond to certain cyclical modes of collaboration between individuals, despite the fact that these hidden modes (almost always adversely) affect the state dynamics. Additionally, the model resulted in graph breakage for a significant class of graphs. Both of these issues made manifest the underlying deficiency of the utility model: it was uncoupled from the state dynamics.

Perceived intelligence, a factor inspired by intuitive notions of network behavior, couples the state and network dynamics in a favorable manner and satisfies the requirement of keeping the network locally adaptive. Namely, perceived intelligence significantly reduces the problem of graph breakage by inducing greater degrees of heterogeneity in network dynamics. It also reduces overshoot in the state dynamics by providing positive feedback between the state and network dynamics. Nevertheless, this positive feedback also results in interesting behaviors regarding network overshoot, the effect of undamped responses to changes in intelligence and leadership. Although we recognize that dynamic leader selection is in fact very favorable for creating robust systems, we also realize that swings in graph structure as nodes unnecessarily gravitate between leaders is sub-optimal behavior for reaching consensus. Such behavior elicits the need for damping the response to changes in intelligence.

Adding a recursive estimation scheme to perceived intelligence, intuitively presented as intelligence history, provides this damping factor. It was found that beyond a certain threshold

($\gamma < O(1)$) increasing the length of history (reducing γ) reduces both the magnitude and frequency of oscillations for network structure, as reflected in the oscillations of the "snapshot" values for the convergence speed and robustness. More importantly, history encodes the ability for a network to implicitly "learn" about leaders over time, thereby making dynamic leader selection even more targeted than in the case of intelligence without history. However, we have found that large histories are detrimental to systems that have a single global leader due to the model's inability to cope with the overshoot of leaders.

While creating the current consensus protocol, we have also developed metrics for consensus performance that take into account network costs. Such cost-normalized effective metrics clearly render the models with intelligence as superior to the original model of [5]. Using these metrics, we have been able to predict performance enhancements if the length of a node's history for intelligence estimation varies inversely with its (unweighted) outdegree. It is expected that incorporating this effect into our model will reduce problems with leader overshoot as well.

More generally, an important characteristic of our model is its ability to easily incorporate newer sophistications and improvements. We have already mentioned the need for making γ depend on the nodal outdegree as well as resolving the dynamics of leaders. Other changes could involve introducing heterogeneity into the cost functions for utility. For example, one might consider nodes that have high costs associated with low intelligence, thereby providing an incentive for such nodes to become leaders. Also, one might consider an additional cost associated with following an exogenous signal. Such a configuration would be relevant to systems that require consensus *accuracy*, an aspect of consensus dynamics that we have not yet explored with the current protocol.

In addition to improving the consensus protocol, we should also extend our analysis to general graphs. It is important to begin analysis with canonical graphs such as complete graphs or stars, as such network extremes offer important results regarding the analytical behavior of the protocol. Having performed such analysis, we can now begin testing our model on more realistic systems such as small-world graphs and Erdös-Rényi random graphs. In particular, we expect that statistical analysis over small-world graphs may provide intuition for dealing with the dynamics of local, but not necessarily global, leaders.

Having discovered numerous frontiers to explore in further research, we pause to reflect on the significance of the framework we have developed. The distributed protocol presented here is by no means guaranteed to provide globally optimal performance. It is a heuristic that tackles the problem of decentralized control over adaptive networks, a domain where globally optimal solutions are often analytically intractable or otherwise constrained in applicability to realistic networks. In this sense, we have created a method analogous to PID control in a domain where the analogue to a linear quadratic regulator is difficult to define. Our system requires very few inputs, just those associated with utility and intelligence history, and it is intentionally built upon a very generic utility model: a sigmoidal benefit subtracted by a quadratic cost. Despite this deceptively simple and intuitive formulation, we have discovered extremely rich behavior and promising performance with respect to addressing noisy consensus dynamics. Overall, we believe our distributed consensus protocol has the ability to develop into a mature technique for controlling networked systems. In this way, the implications of expanding upon the model presented here reach across a range of fields and applications.

A Theorems and Derivations

Theorem A.1. (Gershgorin Circle Theorem) *For an $n \times n$ matrix A with elements a_{ij}, define $R_i := \sum_{j \neq i}^{n} |a_{ij}|$ (the sum of the absolute values of the nondiagonal entries in the i-th row). Every eigenvalue lies in at least one of the Gershgorin discs defined by center a_{ii} and radius R_i in the complex plane.*

Proof. (Adapted from [7]) For the eigenvector $\mathbf{x}$ with eigenvalue λ, choose $i \in \{1, 2..., n\}$ such that $|x_i| = \max_j |x_j| > 0$. Since $A\mathbf{x} = \lambda \mathbf{x}$, then

$$\sum_{j \neq i} a_{ij} x_j = \lambda x_i - a_{ii} x_i$$

Dividing by x_i and taking the absolute value yields:

$$|\lambda - a_{ii}| = \left| \frac{\sum_{j \neq i} a_{ij} x_j}{x_i} \right| \leq \sum_{j \neq i} \left| \frac{a_{ij} x_j}{x_i} \right| \leq \sum_{j \neq i} |a_{ij}| = R_i$$

since $|x_j| \leq |x_i|$ for $j \neq i$. $\square$

Corollary A.2. *The eigenvalues λ of a normalized Laplacian matrix ($\hat{L} \in \mathbb{R}^{N \times N}$) satisfy* $\operatorname{Re} \lambda \in [0, 2]$.

Proof. By construction, $\hat{L}_{ii} = 1$ and the Gershgorin discs satisfy $R_i = 1$. Then, by Theorem A.1, $|\lambda - 1| \leq 1$, so $0 \leq \operatorname{Re} \lambda \leq 2$. $\square$

Lemma A.3. *Σ_{ss} is the solution to $\bar{L}\Sigma_{ss} + \Sigma_{ss}\bar{L}^T - 2\sigma^2 I_{N-1} = 0$.*

Proof.

$$\begin{aligned}
\dot{\Sigma}(t) &= E\left[\dot{\mathbf{y}}\mathbf{y}^T + \mathbf{y}\dot{\mathbf{y}}^T\right] \\
&= E\left[\left(-\bar{L}\mathbf{y} + Q\boldsymbol{\xi}\right)\mathbf{y}^T + \mathbf{y}\left(-\bar{L}\mathbf{y} + Q\boldsymbol{\xi}\right)^T\right] \\
&= -\bar{L}\Sigma - \Sigma\bar{L}^T + QE\left[\boldsymbol{\xi}\mathbf{y}^T\right] + E\left[\mathbf{y}\boldsymbol{\xi}^T\right]Q^T
\end{aligned}$$

The time dependence of $\mathbf{y}$ is given by the standard variation of constants formula: $\mathbf{y}(t) =$

$e^{-\bar{L}t}\mathbf{y}(0) + \int_0^t e^{-\bar{L}(t-\tau)}Q\boldsymbol{\xi}d\tau$, which yields the following:

$$\begin{aligned} E\left[\mathbf{y}\boldsymbol{\xi}^T\right] &= e^{-\bar{L}t}QE\left[\mathbf{x}(0)\boldsymbol{\xi}^T(\tau)\right] + \int_0^t e^{-\bar{L}(t-\tau)}QE\left[\boldsymbol{\xi}(t)\boldsymbol{\xi}^T(\tau)\right]d\tau \\ &= \int_0^t e^{-\bar{L}(t-\tau)}Q\eta^2\delta(t-\tau)d\tau \\ &= Q\sigma^2 \\ \therefore \dot{\Sigma}(t) &= -\bar{L}\Sigma - \Sigma\bar{L}^T + 2\eta^2 I_{N-1} \end{aligned}$$

The steady state solution is constant, i.e. $\lim_{t\to\infty}\dot{\Sigma}(t) = 0$, whereby:

$$\bar{L}\Sigma_{ss} + \Sigma_{ss}\bar{L}^T - 2\eta^2 I_{N-1} = 0$$

□

Proposition A.4. *A homogeneous complete BCC at equilibrium is always unstable when $N \geq N_u$, for N_u given by Equation 3.3a.*

Proof. The homogeneous complete BCC is confined to the line $d = \frac{N-1}{2}\sigma$. Then an equilibrium is guaranteed to be unstable as long as this line only crosses the nullcline for $\sigma < \rho$, which implies $\frac{N-1}{2}\rho > d_0(\rho)$. In dimensionless form, this yields the constraint: $\frac{N-1}{2} > \frac{1}{m\sqrt{r}}$. Therefore, we have:

$$N > \frac{2}{m\sqrt{r}} + 1$$

The smallest integer satisfying this inequality (N_u) is found by adding one to the result and then taking the floor:

$$N_u = \left\lfloor \frac{2}{m\sqrt{r}} + 2 \right\rfloor$$

□

Proposition A.5. *A homogeneous complete BCC at equilibrium is at least marginally stable when $N \leq N_s$, for N_s given by Equation 3.3b.*

Proof. The condition $N \leq \frac{2}{m\sqrt{r}} + 1$ is sufficient for the line $d = \frac{N-1}{2}\sigma$ to cross the nullcline at $\sigma \geq \rho$. However, we want the *only* crossing to occur at $\sigma \geq \rho$. If this is not the first crossing, then based on the shape of $d_0(\sigma)$, it must either be the second crossing or the third crossing. In that case, there must be a smaller value of N for which the nullcline and the line $d = \frac{N-1}{2}\sigma$ touch tangentially. This line has the slope C and satisfies the following two conditions: $C\sigma_s = d_0(\sigma_s)$ and $C = d_0'(\sigma_s)$ for some critical value σ_s. This yields the following solution for $s_s = \sigma_s/\rho$ from which C can be calculated:[32]

$$s_s = \frac{5 - \sqrt{9 - 16r}}{8}$$

When $r > \frac{9}{16}$, there is no real solution, so the condition $N \leq \frac{2}{m\sqrt{r}} + 1$ gives the critical value of N below which the *only* crossing occurs at $\sigma \geq \rho$.

When $r = \frac{9}{16}$, the tangency point is also the point at which the two curves intersect because $d_0''(\frac{5}{8}\rho) = 0$. Then $C\sigma$ must stay above the nullcline for $\sigma > \frac{5}{8}\rho$ and the condition $N \leq \frac{2}{m\sqrt{r}} + 1$ is again strong enough for the *only* crossing to occur at $\sigma \geq \rho$.

When $r \leq \frac{9}{16}$, s_s gives the value at which the two curves are tangent. However, the condition $\frac{N-1}{2} < C$ is not enough for the crossing point to occur at $\sigma \geq \rho$. We also need $C \leq \frac{1}{m\sqrt{r}}$. The latter condition becomes algebraically unfriendly, so it is easier to reason in the following manner. We know that if $C \geq \frac{1}{m\sqrt{r}}$, then the condition $N \leq \frac{2}{m\sqrt{r}} + 1$ will necessarily be strong enough to be an upper bound on N. If C is low enough, then $\frac{N-1}{2} < C$ will be the condition we use to give an upper bound on N. Thus, we can easily get an upper bound on N by just taking the minimum of these two conditions. In practice, this requires less work than checking boundaries for the condition $C \leq \frac{1}{m\sqrt{r}}$.

Since the condition $N < 2C + 1$ is a strict inequality, the largest integer that satisfies it can be found by subtracting one and taking the ceiling. This yields the final result (in nondimensional form):

$$N_s = \begin{cases} \min\left(\left\lfloor \frac{2}{m\sqrt{r}} + 1 \right\rfloor, \lceil 2e_0'(s_s) \rceil\right) & \text{if } r < \frac{9}{16} \\ \left\lfloor \frac{2}{m\sqrt{r}} + 1 \right\rfloor & \text{if } r \geq \frac{9}{16} \end{cases}$$

□

Proposition A.6. *A symmetric BCC star is always unstable when $N \geq N_u^+$, for N_u^+ given by Equation 3.4.*

Proof. The symmetric BCC star lies on the line $d = \frac{N-1}{N}\sigma$ at equilibrium. Then an equilibrium is guaranteed to be unstable as long as this line only crosses the nullcline for $\sigma < \rho$, which implies $\frac{N-1}{N}\rho > d_0(\rho)$. In dimensionless form, this yields the constraint: $\frac{N-1}{N} > \frac{1}{m\sqrt{r}}$.

[32]The values $s = \frac{5 \pm \sqrt{9-16r}}{8}$ are the two locations at which saddle-node bifurcations would appear if N were continuous.

Therefore, we have:

$$N\left(1-\frac{1}{m\sqrt{r}}\right)>1$$

If $m\sqrt{r}\leq 1$, there is no finite integer for N which satisfies this inequality, which means that all symmetric BCC stars are at least marginally stable. For $m\sqrt{r}>1$, we can solve the inequality to get $N>\frac{m\sqrt{r}}{m\sqrt{r}-1}$. This yields the result:

$$N_u^+=\begin{cases}\left\lfloor\frac{m\sqrt{r}}{m\sqrt{r}-1}+1\right\rfloor & \text{if } m\sqrt{r}>1\\ \infty & \text{if } m\sqrt{r}\leq 1\end{cases}$$

□

Proposition A.7. *$\mathcal{H}$ is spanned by symmetric zero-sum 3-cycles. For a complete BCC, the basis of $\mathcal{H}$ is given by a subset of $\binom{N-1}{2}$ of these $\binom{N}{3}$ symmetric zero-sum 3-cycles.*

Proof. We start by finding a basis for the hidden space of a complete BCC. Following the form of Proposition 3.2, we use the same ordering for $\mathbf{m}$ and $\mathbf{n}$ and then write the reduced row echelon form of T_{N+1} ($\operatorname{rref} T_{N+1}$) in terms of $\operatorname{rref} T_N$ (for $N\geq 2$):

$$\operatorname{rref} T_{N+1}=\left(\begin{array}{c|c|c} X & Y & Z\end{array}\right)$$

$$X=\left(\begin{array}{c}\operatorname{rref} T_N\\ \hline 0_{N+1,u_N}\end{array}\right),\quad Y=\left(\begin{array}{c}0_{v_N-1,N+1}\\ \hline I_{N+1}\\ \hline \mathbf{0}_{N+1}^T\end{array}\right),\quad Z=\left(\begin{array}{c}K\\ \hline -\mathbf{1}_{N-1}^T\\ \hline I_{N-1}\\ \hline \mathbf{1}_{N-1}^T\\ \hline \mathbf{0}_{N-1}^T\end{array}\right)$$

where $K\in\mathbb{R}^{v_N-1,N-1}$ is a sparse matrix with the following nonzero entries:

$$K_{ij}=\begin{cases}1 & \text{if } i=\binom{j+1}{2}\\ -1 & \text{if } i=\binom{j+1}{2}+j\end{cases}$$

Removing the "free" columns from T_{N+1} and keeping only the "pivot" columns yields the matrix $I_{v_{N+1}-1}$. Furthermore, the free columns that depend on the new node (the columns in Z) have only 6 nonzero entries. The base case of $\operatorname{rref} T_2$ has no free columns since hidden modes only appear for $N\geq 3$, so all free columns must only contain 6 nonzero entries. Furthermore, examining the placement of the nonzero terms in the columns of Z yields hidden mode vectors that correspond to symmetric zero-sum 3-cycles. Therefore, in this ordering of variables, the basis of hidden mode vectors can be represented by linearly independent symmetric zero-sum 3-cycles. An easy way to pick $\binom{N-1}{2}$ linearly independent hidden modes is to pick one node as the primary node and then make groups of 3 by choosing the other 2 from the set of $N-1$ nodes that does not include the primary node.

The extension to any arbitrary graph is simple. Because the hidden space of the complete BCC contains the hidden space of every other graph, it must be the case that symmetric zero-sum 3-cycles span the hidden space of every graph. Whether or not a particular $\mathbf{h}^{c3}$ is realizable in the specific graph is a separate issue. In general, only a specific linear combination of $\mathbf{h}^{c3}$ vectors will actually be realizable for a given graph. For example, a BCC cycle of N nodes has one hidden mode: $\mathbf{h}^{cN}$. $\square$

Theorem A.8. *The speed of convergence for any connected graph has an upper bound such that* $\lambda_{min} \leq \lambda^*_{min}$.

Proof. If this were not the case, then we could have $\sum_i \lambda_i > N$ for the $N-1$ eigenvalues λ_i in $\hat{L}$. This is impossible since since $\operatorname{tr}\hat{L} = N$ by construction. $\square$

Theorem A.9. *The robustness of convergence for any connected graph has a lower bound such that* $H_2 \geq H_2^*$.

Proof. The solution P to the Lyapunov equation $\bar{L}P + P\bar{L}^T = I_{N-1}$ is:

$$P = \int_0^\infty e^{-\bar{L}t} e^{-\bar{L}^T t} dt$$

and $H_2^2 = \operatorname{tr} P$. In [14], it is shown that:

$$\operatorname{tr}\left(e^A e^{A^T}\right) \geq \sum_{i=1}^{k} e^{2\operatorname{Re}\lambda_i}$$

where λ_i is an eigenvalue of a generic $k \times k$ matrix A. Then, we can bound the square of H_2 as [14]:

$$H_2^2 \geq \int_0^\infty \sum_{i=1}^{N-1} e^{-2\operatorname{Re}\lambda_i t} dt = \sum_{i=1}^{N-1} \frac{1}{2\operatorname{Re}\lambda_i}$$

where λ_i now represents an eigenvalue of $\bar{L}$. The minimum of the right hand side of this equation can easily be found using standard methods of constrained optimization (with the constraint that $\operatorname{tr}\bar{L} = N$). Namely, we can recast the problem into the following:

$$\min_{x_i \in (0,\infty)} \sum_i \frac{1}{x_i} \text{ subject to } \sum_i x_i = 1$$

Rather than use a straightforward but arduous method such as using a Lagrange multiplier to find the critical point, we recognize that the problem does not change if we exchange any two of the variables x_i and x_j. Therefore, the critical point must lie at the point of symmetry in which all x_i are equal. Inspection of the objective function shows that it is decreasing in the parameters and blows up as $x_i \to 0$. Then, the critical point must be a global minimum over the constrained domain $x_i \in (0, \infty) | \sum_i x_i = 1$. It is the point at which the constraint hyperplane lies tangent to a level set of the objective function. Transforming back to the original parameters, we have $\lambda_i = \frac{N}{N-1}$, which yields:

$$H_2^2 \geq \sum_{i=1}^{N-1} \frac{N-1}{2N} = \frac{(N-1)^2}{2N} = \sum_{i=1}^{N-1} \lambda^*_{min} = (H_2^*)^2$$

Then we must have $H_2 \geq H_2^*$. □

B Jacobians for Linearized Network Dynamics

The Jacobian matrices for linearized dynamics are shown below. The vectors are ordered as follows: $\mathbf{m} = (a_{12}, a_{13}, a_{21}, a_{23}, a_{31}, a_{32}, a_{33})$ and $\mathbf{n} = (d_1^{out}, d_2^{out}, d_3^{out}, \sigma_{12}, \sigma_{13}, \sigma_{23}, \sigma_{33})$.

$$\frac{\partial f(\mathbf{m})}{\partial \mathbf{m}} = (Q|R)$$

$$Q = \left(\begin{array}{ccc} B''(\sigma_{12}) - 2\mu & -2(N-2)\mu & B''(\sigma_{12}) \\ -2\mu & B''(\sigma_{13}) - 2\mu(N-2) & 0 \\ B''(\sigma_{12}) & 0 & B''(\sigma_{12}) - 2\mu \\ 0 & 0 & -2\mu \\ 0 & B''(\sigma_{13}) & 0 \\ \hline \multicolumn{3}{c}{0_{2,3}} \end{array}\right)$$

$$R = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & B''(\sigma_{13}) & 0 & 0 \\ -2(N-2)\mu & 0 & 0 & 0 \\ B''(\sigma_{23}) - 2(N-2)\mu & 0 & B''(\sigma_{23}) & 0 \\ 0 & B''(\sigma_{13}) - 2\mu & -2\mu & -2(N-2)\mu \\ B''(\sigma_{23}) & -2\mu & B''(\sigma_{23}) - 2\mu & -2(N-3)\mu \\ 0 & -2\mu & -2\mu & 2B''(\sigma_{33}) - 2(N-3)\mu \end{pmatrix}$$

$$\frac{\partial Tf(\mathbf{n})}{\partial \mathbf{n}} = (S|T)$$

$$S = \left(\begin{array}{ccc} \multicolumn{3}{c}{-2(N-1)\mu I_3} \\ \hline -2\mu & -2\mu & 0 \\ -2\mu & 0 & -2\mu \\ 0 & -2\mu & -2\mu \\ 0 & 0 & -4\mu \end{array}\right)$$

$$T = \left(\begin{array}{cccc} B''(\sigma_{12}) & (N-2)B''(\sigma_{13}) & 0 & 0 \\ B''(\sigma_{12}) & 0 & (N-2)B''(\sigma_{23}) & 0 \\ 0 & B''(\sigma_{13}) & B''(\sigma_{23}) & (N-3)B''(\sigma_{33}) \\ \hline \multicolumn{4}{c}{2\left(B''(\sigma_{12}), B''(\sigma_{13}), B''(\sigma_{23}), B''(\sigma_{33})\right) I_4} \end{array}\right)$$

References

[1] Agaev, R., and Chebotarev, P. On the spectra of nonsymmetric laplacian matrices. *Linear Algebra and Its Applications 399* (2005), 157–168.

[2] Baxter, L. A. Trajectories of relationship disengagement. *Journal of Social and Personal Relationships 1*, 1 (1984), 29–48.

[3] Bishop, C., et al. *Pattern recognition and machine learning*, vol. 4. springer New York, 2006.

[4] Buckingham, E. On physically similar systems. *Phys. Rev. 4* (1914), 345–76.

[5] Do, A., Rudolf, L., and Gross, T. Patterns of cooperation: fairness and coordination in networks of interacting agents. *New Journal of Physics 12*, 6 (2010), 063023.

[6] Heckathorn, D. D. The dynamics and dilemmas of collective action. *American Sociological Review* (1996), 250–277.

[7] Horn, R. A., and Johnson, C. R. *Matrix analysis*. Cambridge university press, 1990.

[8] Leonard, N., and Olshevsky, A. Cooperative learning in multi-agent systems from intermittent measurements. *arXiv preprint arXiv:1209.2194* (2012).

[9] Nicholson, W., and Snyder, C. *Microeconomic theory: basic principles and extensions*. South-Western Pub, 2011.

[10] Oliver, P., Marwell, G., and Teixeira, R. A theory of the critical mass. i. interdependence, group heterogeneity, and the production of collective action. *American Journal of Sociology* (1985), 522–556.

[11] Strogatz, S. Nonlinear dynamics and chaos: with applications to physics, biology, chemistry and engineering.

[12] Strogatz, S. H. Love affairs and differential equations. *Harmony of the World: 75 Years of Mathematics Magazine 61* (2007), 211.

[13] Weber, R. On the gittins index for multiarmed bandits. *The Annals of Applied Probability* (1992), 1024–1033.

[14] Young, G. Personal communication. *10 April 2013*.

[15] Young, G., Scardovi, L., Cavagna, A., Giardina, I., and Leonard, N. Starling flock networks manage uncertainty in consensus at low cost. *PLOS Computational Biology 9*, 1 (2013), e1002894.

[16] Young, G., Scardovi, L., and Leonard, N. Robustness of noisy consensus dynamics with directed communication. In *American Control Conference (ACC), 2010* (2010), IEEE, pp. 6312–6317.

[17] Young, G., Scardovi, L., and Leonard, N. Rearranging trees for robust consensus. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on* (2011), IEEE, pp. 1000–1005.

This paper represents my own work in accordance with University regulations.