Uncertain Network Dynamics

Noah Williams*
Miami Herbert Business School
University of Miami
noahwilliams@miami.edu

February 14, 2026

Abstract

I study the dynamics of flows on networks under uncertainty. I analyze the dynamic transmission across a network of shocks with uncertain dynamic properties, which can capture both exogenous stochastic shocks and uncertainty in the network structure. I show how to map uncertain network dynamics into a control system, which allows me to apply tools from robust control theory to analyze system performance. I study two main classes of applications: production networks and consensus problems with local information diffusion. I characterize the response of these networks to perturbations, whether due to (1) stochastic shocks with specified properties, where performance is measured by the volatility of the network output, or (2) uncertain but norm-bounded inputs, leading to a robust measure of worst-case shock amplification. I analyze several examples to illustrate how the network topology affects shock propagation and amplification.

1 Introduction

1.1 Statics to Dynamics on Networks

Decentralized, networked interactions play an increasingly central role in modern economic systems. From production and supply chains to financial linkages and informational flows, many key aggregate phenomena emerge from interactions across agents, firms, or sectors with particular, localized or specialized links. As a result, a growing literature has sought

*I thank Ben Golub, Matt Jackson, and Chao Wei for helpful comments and discussion on the topics of this paper.

to understand how network structure shapes aggregate outcomes, and the resilience and fragility of such systems in their response to local disturbances, failures, or shocks.[1]

While the literature on economic networks has generated valuable insights, most of it focuses on static models, with a few notable exceptions discussed below. The network is typically treated as a fixed mapping from shocks to outcomes, without modeling the dynamics of transmission over time or the uncertainty in either the shocks or the network structure. Such models are useful in characterizing the long-run changes in outcomes as a shock propagates cross-sectionally through a network, but cannot address the dynamic propagation and interaction of shocks.

By contrast, I focus on the dynamics of flows on networks. I introduce a framework that explicitly incorporates both dynamic propagation and uncertainty, using tools from control theory. I consider a continuous-time model in which the system evolves according to a differential equation on a directed, weighted graph. The network is represented by an asymmetric matrix, capturing both the directionality and heterogeneity of links. The resulting system describes diffusion of shocks or perturbations over time and through the network, which I analyze using methods from control theory.

In particular, I develop a state-space model of perturbed Laplacian dynamics on networked systems. Related Laplacian dynamics are a common method of studying diffusion on networks.[2] I focus on two general classes of models which have a Laplacian structure: production networks and information propagation. Each of these classes has seen wide applications in economics and related fields.

A main motivating example is the classic model of Long and Plosser (1983), which introduced a dynamic production network linking sectoral outputs through intertemporal supplier relationships. In that model, each sector's output follows a stochastic difference equation whose innovations are passed through to downstream sectors via a fixed input-output structure. Static or steady-state versions of this model have been analyzed in the networks literature, following Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012). Carvalho and Tahbaz-Salehi (2019) provide a good summary and overview.[3] These papers analyze amplification effects in equilibrium, using static models of sectoral interactions.

While many production network models are static, the original Long and Plosser (1983) model has a dynamic structure which has been utilized in the business cycle literature.[4] Here

[1]The economic literature on networks includes important and comprehensive surveys by Jackson (2008), Pass (2019), and Goyal (2023), which study the formation and economic implications of networked interactions. Each of these books contain detailed references to the literature.

[2]See Newman (2018), Veerman and Lyons (2020), and Van Mieghem (2023).

[3]Other contributions emphasizing the role of input-output networks and sectoral linkages include Baqaee and Farhi (2019), Baqaee and Farhi (2020). Empirical and international applications include Boehm, Flaaen, and Pandalai-Nayar (2019), and Huo, Levchenko, and Pandalai-Nayar (2025).

[4]Important later work includes Horvath (1998), Dupor (1999), Gabaix (2011) and Atalay (2017), who

I show that the equilibrium dynamics can be mapped into a version of a graph Laplacian.

Some recent work also analyzes related dynamic network models. In a model close to mine, Liu and Tsyvinski (2024) study a production network with dynamics driven by adjustment costs. They analyze the deterministic dynamics of the network economy as it recovers from a temporary negative productivity shock, and show that temporary shocks have disproportionate welfare effects on upstream sectors. Schaal and Taschereau-Dumouchel (2025) study a dynamic production network with heterogeneous time-to-build, generalizing the model I use. Like me, they use spectral and frequency-domain methods to show how cycles in the network can generate endogenous fluctuations. Carvalho, Covarrubias, and Nuño (2025) study a nonlinear dynamic network model under uncertainty, and show that the endogenous response of capital amplifies the welfare cost of business cycles.

While I study the impact of shocks in a fixed network, a related literature studies how the network itself may respond to shocks and lead to amplification. This includes work studying how uncertainty and incentives can affect the formation and fragility of networks.[5] Related papers show how shocks can reconfigure network links and generate cascades.[6] Complementary research by Elliott and Jackson (2026) studies how, in an existing network, reallocation following disruptions can amplify short-run responses to shocks. My results can be interpreted as a stress test for networks that are either technologically fixed or generated endogenously in these models.

My other main application considers local information transmission over a network. There is a broad literature on learning in social networks, surveyed by Acemoglu and Ozdaglar (2011) and Golub and Sadler (2016), which studies conditions under which distributed and local updating can lead to convergence, mis-learning, or persistent disagreement.[7] While there are a variety of Bayesian and non-Bayesian learning approaches, a number of recent papers have focused on the DeGroot (1974) approach to learning, as in DeMarzo, Vayanos, and Zwiebel (2003) and Golub and Jackson (2010).[8]

highlight the potential for idiosyncratic shocks to generate aggregate fluctuations. Foerster, Sarte, and Watson (2011) study the implications of a related model for the factor structure of shocks, and vom Lehn and Winberry (2022) add an investment network along with the production network.

[5]See, for example, Acemoglu and Azar (2020), Elliott, Golub, and Leduc (2022), and Kopytov, Mishra, Nimark, and Taschereau-Dumouchel (2024).

[6]See Baqaee (2018) and Taschereau-Dumouchel (2025).

[7]Early contributions on informational cascades by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), and word-of-mouth communication by Ellison and Fudenberg (1993) and Ellison and Fudenberg (1995) showed that even simple networked interactions can lead to herding and inefficiency. Later contributions include Bala and Goyal (1998), Gale and Kariv (2003), and Acemoglu, Dahleh, Lobel, and Ozdaglar (2011).

[8]Further applications and modifications of this approach include Ghosh (2021), Cheng, Han, Tong, Wu, and Xing (2023), Banerjee and Compte (2024), and Amir, Arieli, Ashkenazi-Golan, and Peretz (2025).

1.2 Contribution of the Paper

My main approach is to map network dynamics into a control system, making use of "Laplace" methods in two important ways. First, the dynamics on the network are governed by a differential equation involving the graph Laplacian, which characterizes how shocks or other inputs diffuse or propagate across the nodes of the network. The Laplacian essentially becomes the state matrix in a control system. Second, I take the Laplace transform of the system, converting time-domain behavior into the frequency domain. This double use of Laplace techniques allows me to reframe the perturbed or stochastic network dynamics as a transfer function which maps inputs to outputs. I then apply methods from control theory to characterize the volatility of network outputs and the robustness or fragility of different network structures.[9]

The control system approach allows for two broad frameworks to study the performance of the network under uncertainty. In the first, I analyze the impact of stochastic shocks on the network output. In this case, the input is a process with specified properties.[10] This naturally leads to measures of network performance which are tied to the variance of the system output, a measure which is known as the $\mathcal{H}_2$ norm. In the frequency domain, this measure averages network performance across frequencies.

But the control system approach also accommodates a broader notion of uncertainty, when the input is an uncertain process, bounded in a norm, that is potentially correlated with the state vector. This approach, which is common in robust control, covers general exogenous shocks processes with unspecified correlations, as well as broader uncertainty about the structure and dynamics of the network itself. Having only a norm bound leads to worst-case analysis, which provides a robust bound on the network performance. In the frequency domain, this performance measure considers the worst-case performance across frequencies, a measure which is known as the $\mathcal{H}_\infty$ norm.[11] The worst-case performance analysis reveals how certain combinations of shocks across the network could be very damaging to the system performance.

Throughout, I treat the network structure as fixed, analyzing the properties of a given network subjected to shocks of uncertain structure and timing. In the terminology of control theory, I focus here on analysis not synthesis. The methods I develop allow me to assess both the performance (how well the network transmits or attenuates inputs) and robustness (how

[9] The control-theoretic framework employed here is developed in Zhou, Doyle, and Glover (1996), Zhou and Doyle (1998) and Skogestad and Postlethwaite (2005). Hansen and Sargent (2008) pioneered the development and application of robust control methods in economics.

[10] In the baseline case, the shocks are Gaussian white noise. Correlated shocks can be analyzed by expanding the state space.

[11] Related frequency domain analysis of robustness in economics includes Onatski and Williams (2003), Brock and Durlauf (2005), and Brock, Durlauf, and Rondina (2013).

sensitive it is to model uncertainty or disturbance variation) of different network structures. In particular, I identify network architectures that perform well in the face of uncertain disturbances, and diagnose which topologies are fragile or prone to amplification. These results are useful on their own, and provide key inputs for decisions on routing or allocations across the network, as well as analysis of network formation and adaptation.

The theoretical literature on networks has developed measures of the importance of nodes, typically summarized through centrality indices. A large class of such measures arises from spectral or eigenvalue decompositions of the network matrix, which in turn are tied to long-run dynamic properties.[12]

I show that for measures of network performance under uncertainty, what matters instead are the singular values of the control system transfer function. The singular values govern how shocks are amplified, and the $\mathcal{H}_2$ performance metric can be written as the sum of squared singular values integrated across frequencies, while the $\mathcal{H}_\infty$ norm is the maximal singular value across frequencies. When networks are undirected and hence have symmetric adjacency matrices, the singular values are the absolute value of the eigenvalues. But for directed or asymmetric networks, the singular values capture the relevant responses.

2 Two Motivating Examples

2.1 Production Networks

To motivate the framework developed in this paper, I begin with the dynamic input-output model introduced by Long and Plosser (1983), which provides a classic example of an economy where sectoral interconnections give rise to aggregate dynamics.

The economy consists of N sectors, each producing a distinct perishable good that can be used either as an intermediate input in production or consumed directly. Each sector is operated by a representative firm that combines labor and intermediate inputs from other sectors to produce output. The production function for firm i in period $t+1$ is given by

$$Y_{i,t+1} = Z_{i,t+1} L_{i,t}^{b_i} \prod_{j=1}^{N} X_{ij,t}^{a_{ij}},$$

where $L_{i,t}$ denotes labor employed by firm i, $X_{ij,t}$ is the quantity of good j used by firm i as an input, and $Z_{i,t+1}$ is a stochastic productivity shock. The exponents b_i and a_{ij} are production elasticities, with the assumption of constant returns to scale implying $b_i + \sum_j a_{ij} = 1$. Here

[12] Eigenvector and Katz-Bonacich centrality, from Katz (1953) and Bonacich (1987), are used in network games following Ballester, Calvó-Armengol, and Zenou (2006), as well as the production network and information diffusion literature discussed above. See Bloch, Jackson, and Tebaldi (2023) for an overview of network centrality measures.

b_i is the labor share of firm i, so $\sum_j a_{ij} = 1 - b_i \equiv \beta_i$ is the input share. Input shares govern the speed of propagation through the network, and play a key role in the analysis to follow.

On the demand side, a representative consumer chooses labor supply and consumption to maximize utility. Preferences are separable and logarithmic over leisure ℓ_t and the vector of consumptions $C_t = \{C_{i,t}\}$ across goods:

$$U(C_t, \ell_t) = \theta_0 \log \ell_t + \sum_{i=1}^{N} \theta_i \log C_{i,t},$$

where the weights θ_0 and θ_i capture the importance of leisure and each good in utility. The time endowment constraint requires that leisure plus total labor supplied across sectors equals a fixed time resource h, i.e.,

$$\ell_t + \sum_{i=1}^{N} L_{i,t} = h.$$

Markets clear each period: the output of each sector equals the sum of its consumption and use as an intermediate input by other sectors,

$$Y_{j,t} = C_{j,t} + \sum_{i=1}^{N} X_{ij,t}.$$

Under these conditions, equilibrium allocations can be solved explicitly. The allocation of labor across sectors is constant at $L_{i,t} = \bar{L}_i$, and consumers choose a constant level of leisure $\ell_t = \bar{\ell}$. Intermediate inputs are demanded in fixed proportions of output, so that $X_{ij,t} = \bar{x}_{ij} Y_{j,t}$ for some constants $\bar{x}_{ij}$.

Substituting these relationships into the production function and taking logarithms yields a dynamic linear equation for output. In particular, the log of sectoral output evolves according to

$$\log Y_{i,t+1} = k_i + \sum_{j=1}^{N} a_{ij} \log Y_{j,t} + \log Z_{i,t+1},$$

where k_i is a constant and a_{ij} is the share of good j used by sector i in production.

Vectorizing this system across sectors, letting $y_t = \log Y_t - \log \bar{Y}$ denote log deviations from steady state, and taking a continuous-time limit, we can write the dynamics as:

$$\dot{y} = -(\mathrm{I} - A)y + z \equiv -Ly + z,$$

where A is the matrix of input-output coefficients, z collects the productivity shocks. Here $L = \mathrm{I} - A$ is a Laplacian operator that governs the flow of shocks through the production network. If there were no labor input so $b_i \equiv 0$, then constant returns to scale would make

the matrix A row-stochastic with unit input shares. In this case, the operator L is a random-walk graph Laplacian, a key object in the analysis of dynamic processes on networks. But here I assume that each sector uses labor, so intermediate shares satisfy $\beta_i < 1$. Thus the adjacency A is sub-stochastic and $-L$ is a Hurwitz matrix, so the production network is dynamically stable. The Laplacian dynamics propagate the sectoral shocks over time across the network, but their effects decay due to leakage into labor.

Much of the subsequent literature, as summarized in Carvalho and Tahbaz-Salehi (2019), has adopted a static version of this model to study the long-run cross-sectional diffusion of shocks. In the static setting, output is determined by the fixed-point equation

$$y = Ay + z \quad \Rightarrow \quad y = (\mathrm{I} - A)^{-1} z,$$

where $(\mathrm{I} - A)^{-1}$ is the so-called Leontief inverse. This matrix captures the full set of direct and indirect input relationships among sectors and admits the power series expansion

$$(\mathrm{I} - A)^{-1} = I + A + A^2 + A^3 + \dots,$$

in which the (i, j) entry measures the cumulative influence of sector j on sector i through all input chains of any length.

Here I incorporate dynamics both through output that evolves according to the differential equation above, and through dynamic shocks $z(t)$ that may evolve over time, possibly stochastically or with persistent autocorrelation. I develop an approach to analyze uncertainty in both the shocks themselves and in the network structure, providing a general treatment of uncertainty and shock transmission in networked economies.

The production economy can be interpreted as a weighted, directed graph. Each node represents a sector, and a directed edge from node j to node i exists if sector j supplies intermediate inputs to sector i. The edge weight a_{ij} reflects the share of input j used by firm i, and the direction reflects the flow of goods from suppliers to users. Because these flows are inherently directional, the corresponding adjacency matrix A is generally asymmetric. Below I develop some analytic results in the special cases where networks are symmetric, or where all industries have a common input share parameter $\beta_i = \beta < 1$ for all i.

2.2 Information Propagation and Consensus

The dynamic input-output structure discussed above naturally generalizes to a broader class of systems governed by network-based propagation and diffusion dynamics. Let the underlying network be represented by an adjacency matrix A, where $a_{ij} > 0$ denotes a directed edge from node j to node i. The associated in-degree matrix is $D = \mathrm{diag}[\sum_j a_{ij}]$, and the graph Laplacian is defined as $\mathcal{L} = D - A$. As described by Veerman and Lyons (2020), this

Laplacian operator governs the flow of quantities across the network and forms the foundation for two closely related dynamic processes: consensus and diffusion. But for comparison with the production case, rather than this (un-normalized) Laplacian $\mathcal{L}$, I focus on a random walk Laplacian $L = \mathrm{I} - D^{-1}A$. Equivalently, I restrict attention to adjacency matrices which are row-stochastic, $\sum_j a_{ij} = 1$, so $D = I$.

In the consensus formulation, the vector of node states $x(t)$ evolves according to the differential equation:

$$\dot{x} = -Lx,$$

where x is a column vector. This process models the gradual alignment of node values over time, and has been widely studied in models of belief formation, social learning, synchronization, and opinion dynamics.

To make these dynamics more concrete, consider the example of DeGroot (1974) learning in a social network.[13] Each node in the network represents an agent, and the state $x_i(t)$ corresponds to agent i's belief at time t about some fixed underlying state of the world. Agents update their beliefs continuously by averaging with their neighbors. If agent i places weight a_{ij} on agent j's belief, then the dynamics become:

$$\dot{x}_i = \sum_j a_{ij}(x_j - x_i) = -\sum_j L_{ij}x_j,$$

which is exactly the consensus process described above.

This framework captures the gradual convergence of beliefs in a network. As time evolves, all agents' beliefs tend to align, eventually reaching consensus. Beyond their importance in economics, such consensus models are widely used for multi-agent control systems where synchronization is an objective, as described in Mahmoud, Oyedeji, and Xia (2021). Consensus dynamics lead to uniform (consensus) beliefs in the long-run, where the system does not converge to zero, but to a nontrivial constant vector determined by the initial state. If new information is introduced via a perturbation, eventually it transmits through the whole network and is learned by everyone. This behavior corresponds to the presence of a zero eigenvalue in the Laplacian matrix, with a constant vector as the associated eigenvector.

The presence of a zero eigenvalue implies that the consensus dynamics are only neutrally stable. While this is appealing as a model of learning, it raises some difficulties in analyzing responses to shocks. If we introduce even small shocks, the system will diverge and we cannot meaningfully analyze the long-run performance of the network.

One resolution is to study not the levels of beliefs, but their differences across agents. In the DeGroot example, the disagreement between agents can be expressed in terms of belief differences, which evolve according to stable dynamics. In particular, letting 1_N denote the

[13] See also Jackson (2008) and Golub and Jackson (2010).

N-dimensional vector of ones, we set:

$$C = \check{C} \equiv \mathrm{I} - \frac{1}{N} 1_N 1_N^\top, \tag{1}$$

so that $y_t = \check{C} x_t$ tracks the difference of beliefs from the average, $y_{i,t} = x_{i,t} - \frac{1}{N}\sum_{j=1}^{N} x_{j,t}$. Note that by construction $\check{C} 1_N = 0$. Since the nullspace of the Laplacian L consists of constant vectors, projecting the state onto the orthogonal complement (by eliminating the mean) yields a system that is asymptotically stable. This corresponds to tracking the evolution of disagreement, which always decays to zero. In this paper I focus on such disagreement dynamics, while Appendix A.1 discusses an alternative with dampened consensus dynamics.

3 Perturbed Network Dynamics

I now turn to the analysis, introducing methods from control systems to study perturbed network dynamics.

3.1 Setup

I now deal with a general network, represented by an $N \times N$ adjacency matrix A, and Laplacian $L = I - A$. I suppose that we observe or care about a linear combination of the state, represented by the output $y = Cx$, where C may select specific nodes or aggregate across the network.

Then the general perturbed dynamic system is given by:

$$\begin{aligned} \dot{x} &= Mx + Bu, \\ y &= Cx, \end{aligned} \tag{2}$$

where $x \in \mathbb{R}^N$ is the state vector, u is an input or disturbance, $M = -L = A - \mathrm{I}$, and B maps the input into the state space. This defines a continuous-time, linear, multi-input, multi-output (MIMO) control system.[14]

In the production model, the productivity shocks directly enter as exogenous stochastic inputs in this way. In the consensus model, this allows for the inclusion of exogenous news or information that affects some agents directly. For instance, if a subset of agents receives signals, the system's evolution is governed by the interplay between internal averaging (consensus) and external input (shocks). I also describe below how $u(t)$ may represent uncertainty in the network or its dynamics. Mapping the network dynamics into a standard control system allows me to use robust control methods to evaluate the stability and performance of network dynamics in the face of uncertainty and shocks.

[14] See Skogestad and Postlethwaite (2005) for more background and details on such control systems. While the continuous time approach is slightly simpler, the results all have parallels in discrete time.

3.2 Laplace Transform the Laplacian

Having established the network dynamics in the time domain, I now turn to a frequency-domain analysis using tools from control theory. This approach allows for a flexible treatment of how shocks and uncertainties of different types, frequencies, and sources propagate through a networked system. Unlike traditional static models, which characterize only the long-run cross-sectional response to shocks, this framework analyzes how the structure of the network affects both the timing and magnitude of shock transmission across the network.

The Laplace transform is a classical tool that converts a time-domain system into the frequency domain. For a complex variable $s = \sigma + j\omega$ with $\sigma, \omega \in \mathbb{R}$ and $j = \sqrt{-1}$, the Laplace transform of a function $f : \mathbb{R}_+ \to \mathbb{R}$ is defined as

$$F(s) = \int_0^\infty e^{-st} f(t)\, dt.$$

Applying this to the linear control system $\dot{x} = Mx + Bu$, with initial condition $x(0) = 0$, we obtain the transformed system:

$$\begin{aligned} sX(s) &= MX(s) + BU(s), \\ Y(s) &= CX(s), \end{aligned}$$

where $X(s)$, $U(s)$, and $Y(s)$ are the Laplace transforms of the state, input, and output vectors. Solving algebraically yields the transfer function:

$$Y(s) = C(sI - M)^{-1}BU(s) \equiv G(s)U(s),$$

which maps frequency-domain inputs $U(s)$ to outputs $Y(s)$.

Importantly, the transfer function is the Laplace transform of the familiar impulse response functions. Specifically, the matrix function

$$g(t) = Ce^{Mt}B$$

is the time-domain impulse response: it gives the output of the system at time t in response to a unit impulse input at time 0. Then the transfer function is

$$G(s) = \int_0^\infty e^{-st} g(t)\, dt.$$

Thus, the transfer function is a frequency-domain summary of the impulse response, and characterizes the system's dynamic response to shocks of all temporal shapes.

3.3 Static Limit: The Leontief Inverse

One advantage of this framework is that it nests the familiar static analysis as a special case. Suppose we take the frequency zero limit: $s = 0$. Then the transfer function becomes

$$G(0) = C(-M)^{-1}B.$$

Now use suppose $B = C = I$, which means that we have a full vector of inputs or shocks that can affect each state variable, and that the output of interest is the full state vector. Then, using the definition of M, we have:

$$G(0) = (I - A)^{-1},$$

which is exactly the Leontief inverse discussed above. This captures the long-run, steady-state response of the system to a permanent unit shock, assuming linear propagation through the network. This shows that traditional input-output models correspond to evaluating the transfer function at frequency zero and using identity output and input mappings.

The low frequency behavior is important, and I show below that some key properties of a class of models, which includes the production network model, are governed by their behavior at frequency zero. However in dynamic settings, we are generally interested in the dynamic response of shocks, which may themselves have cross-sectional and temporal correlations. Thus to characterize the performance of a dynamic model in practical settings we need to analyze the response across the frequency spectrum. The transfer function $G(j\omega)$ captures the sensitivity of outputs to shocks at each frequency ω, allowing us to analyze how different types of shocks affect the system over time.

3.4 The Scalar Case

To build further intuition, consider the scalar case where $x(t)$ and $u(t)$ are real-valued and $G(s)$ is a scalar function. Then:

$$Y(s) = G(s)U(s), \quad \text{so} \quad y(t) = g * u(t),$$

where $*$ denotes convolution. In this case, the gain of the system at frequency ω is simply the modulus of the transfer function:

$$|G(j\omega)| = \frac{|Y(j\omega)|}{|U(j\omega)|}.$$

The gain measures how the system amplifies, if $|G(j\omega)| > 1$, or attenuates, if $|G(j\omega)| < 1$, inputs at frequency ω.

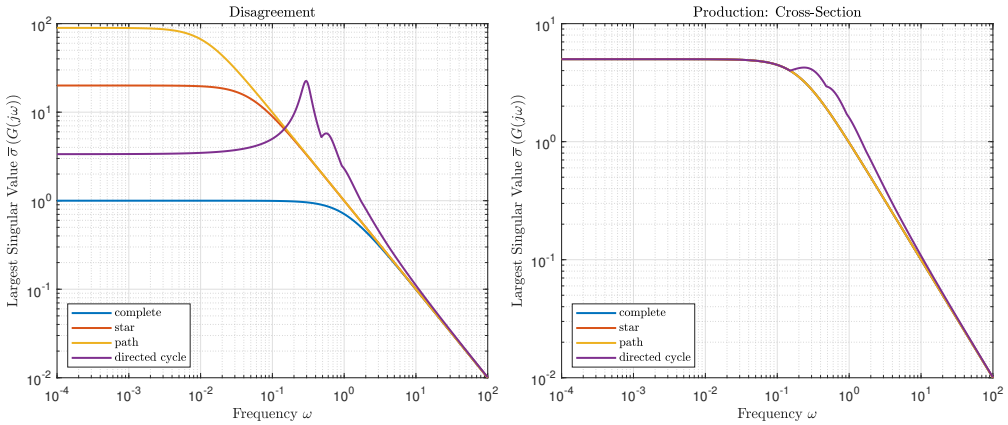


Figure 1: Largest singular value $\overline{\sigma}(G(j\omega))$ for several example networks for disagreement dynamics (left panel) and production networks (right panel).

3.5 MIMO Systems and Singular Values

In the remainder of the paper I will follow the standard convention in control theory and use the same (lower-case) notation for both a variable in the time domain and its Laplace transform in the frequency domain.

In multi-input, multi-output (MIMO) systems, the transfer function $G(j\omega)$ is a matrix, and both input and output are vectors. To assess system performance, we study how $G(j\omega)$ transforms input vectors into outputs. The appropriate multidimensional generalization of gain is the largest singular value:

$$\overline{\sigma}(G(j\omega)) = \max_{u(\omega)\neq 0} \frac{\|G(j\omega)u(\omega)\|_2}{\|u(\omega)\|_2}.$$

This gives the maximum amplification at frequency ω, over all possible input directions $u(\omega)$.

Figure 1 shows the maximal singular value $\overline{\sigma}(G(j\omega))$ as a function frequency for several example networks that I study in detail in Section 5.1 below. The left panel shows the results for disagreement dynamics as in Section 2.2, while the right panel shows simple production networks as in Section 2.1.[15] For reasons made clear below, many of the networks have singular values that peak at frequency zero, with the exception being disagreement dynamics on the directed cycle. Furthermore, embedding different network topologies in a production network with homogeneous labor shares eliminates most of the differences, at least in terms of their largest singular values.

Recall that the singular values of a complex matrix G are always well-defined, and are given by the square roots of the eigenvalues of $G^H G$, where $G^H = \left(\overline{G}\right)^T$ is the Hermitian or

[15]Specifically, the plot uses networks of size $N = 21$ with homogeneous input share parameter $\beta = 0.8$ in the production case.

conjugate transpose:

$$\sigma_i(G) = \sqrt{\lambda_i(G^H G)}.$$

Recall as well that a matrix is normal if $G^H G = G G^H$. Symmetric matrices, which govern the dynamics of undirected networks, are an important special case of normal matrices. Normal matrices also include other examples, such as rotation matrices, which govern the dynamics of cyclical networks. For normal matrices, the singular values equal the absolute values of the eigenvalues, $\sigma_i(G) = |\lambda_i(G)|$. But for general asymmetric matrices, such as for general directed networks, there is no clear relation between the eigenvalues and singular values.

Parallel to eigenvectors, the left and right singular vectors are useful to diagnose input and output directions. The largest singular value of a matrix G satisfies:

$$\bar{\sigma}(G) = \max_{\|v\|_2=1} \|Gv\|_2 = \|G\|_2,$$

where the last equality defines the induced matrix 2-norm. A unit vector $\bar{v}$ achieving this maximum is a right singular vector, which solves $G^H G\bar{v} = \bar{\sigma}^2\bar{v}$. The corresponding left singular vector is $\bar{u} = G\bar{v}/\bar{\sigma}$, equivalently $GG^H\bar{u} = \bar{\sigma}^2\,\bar{u}$.

Most network analyses in economics focus on the eigenvalues and eigenvectors of the adjacency matrix A. Measures such as eigenvector centrality or Bonacich centrality depend on the adjacency matrix A, and need not capture dynamic amplification in the presence of asymmetry or non-normality. A normal matrix has orthogonal eigenvectors, which thus form a basis, and the eigenvalues capture how the system evolves. But when a matrix is not normal, its eigenvectors need not be orthogonal. The eigenvalues still measure how the system evolves along its eigenvectors, but in response to perturbations or inputs the system will rarely evolve in those directions. Any disturbance will activate distinct, interacting modes, and may produce amplification not captured by the eigenvalues.

An example adapted from Skogestad and Postlethwaite (2005) illustrates the point. Consider a two-node directed network with a directed edge from node 2 to node 1 having a strong weight of 100, and no edge from node 1 to node 2. Suppose as well that node 2 is subject to potential shocks, but that we care about the output from node 1. We can represent this as a dynamic control system (2) with:

$$M = \begin{bmatrix} -1 & 100 \\ 0 & -1 \end{bmatrix}, \qquad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \qquad C = \begin{bmatrix} 1 & 0 \end{bmatrix}. \tag{3}$$

The system matrix M has a double eigenvalue at -1, which by itself suggests a uniformly stable system with no dramatic amplification. However, M is highly non-normal. With the double eigenvalue, it has only a single eigenvector, proportional to the unit vector $[1, 0]^\top$, and so clearly does not form a basis. In particular, the eigenvalues don't capture movement in the second dimension, induced by the perturbations at the second node. For example,

the response to a unit impulse at time zero is $g(t) = 100t\exp(-t)$ which peaks at $g(1) = 100\exp(-1) \approx 36.8$, showing significant amplification. The eigenvalues don't reflect the gain from input, a shock to node 2, on the output of node 1.

While eigenvalues describe how the internal state of a system evolves, singular values describe how the system transforms inputs into outputs, and hence here how shocks propagate. In some settings, with dynamics captured by a normal matrix, we are able to reduce the analysis to eigenvalue properties. But in general, singular vectors capture the interactions across modes, which may amplify or dampen the system's response to shock. Thus singular values are key objects of the analysis.

3.6 System Norms

While the singular value at each frequency gives the maximum gain, to summarize the system's overall performance across frequencies, we compute norms of the transfer function. I focus on the two standard measures of system performance: the $\mathcal{H}_\infty$ and $\mathcal{H}_2$ norms.

In particular, the $\mathcal{H}_\infty$ norm measures the worst-case amplification from inputs to outputs across all directions and frequencies:

$$\|G(s)\|_\infty = \sup_\omega \overline{\sigma}(G(j\omega)) = \sup_\omega \max_{u(\omega)\neq 0} \frac{\|G(j\omega)u(\omega)\|_2}{\|u(\omega)\|_2} = \sup_\omega \|G(j\omega)\|_2 \tag{4}$$

This is the key object in robust control, and corresponds to the largest possible gain the system can produce, given any bounded input. This norm is especially relevant when the properties of the shocks are unknown. Rather than specifying the stochastic properties of the input $u(t)$, we take it as an unknown perturbation that can either represent a true shock of uncertain structure or a modeling error. The $\mathcal{H}_\infty$ norm bounds the performance of the network in the presence of these uncertain shocks.

Alternatively, the $\mathcal{H}_2$ norm measures average performance across frequencies:

$$\|G(s)\|_2 = \sqrt{\frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{tr}(G(j\omega)^H G(j\omega))\, d\omega} = \sqrt{\frac{1}{2\pi}\int_{-\infty}^{\infty} \sum_i \sigma_i^2(G(j\omega))\, d\omega}. \tag{5}$$

If $u(t)$ is a Gaussian white noise process and the state matrix M is Hurwitz, then $\|G(s)\|_2 = \sqrt{\mathbb{E}[y(t)^H y(t)]}$ represents the steady-state root mean square magnitude of the output. While the $\mathcal{H}_\infty$ norm depends on the maximum singular value, the second expression above shows that all of the singular values, averaged across frequencies, contribute to the $\mathcal{H}_2$ norm.

For single-input single-output systems, the transfer function is a scalar. Thus the $\mathcal{H}_\infty$ norm is just the maximal gain, the peak of the frequency response $|G(j\omega)|$ across frequencies. On the other hand, the $\mathcal{H}_2$ norm is proportional to the area under the square of the frequency

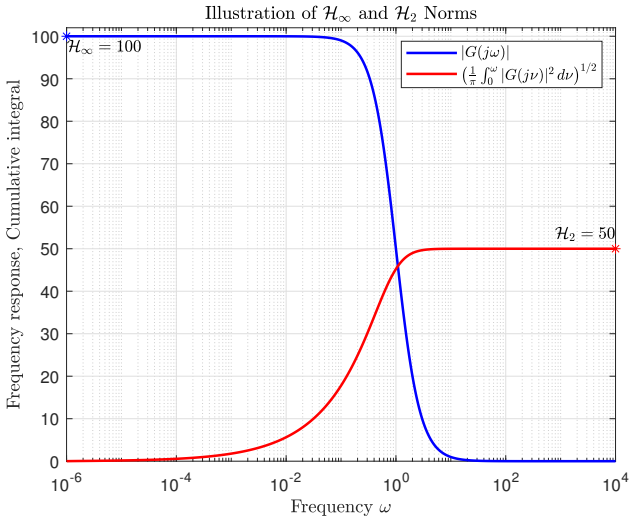


Figure 2: Illustration of the $\mathcal{H}_\infty$ and $\mathcal{H}_2$ norms in an example. Shown are the magnitude of the transfer function and square root of the cumulative area under its square.

response $|G(j\omega)|^2$. These norms are illustrated in Figure 2, for the example (3) from the previous section. The transfer function from the scalar input to the scalar output is:

$$G(s) = C(sI - M)^{-1}B = \frac{100}{(s+1)^2}.$$

Thus the $\mathcal{H}_\infty$ norm of this transfer function is

$$\|G\|_\infty = \sup_{\omega\in\mathbb{R}} |G(j\omega)| = \sup_{\omega} \frac{100}{\omega^2+1} = 100,$$

so the worst-case input-output amplification of energy-bounded disturbances is 100. This is shown in the figure as the peak at frequency zero. On the other had, we have:

$$\|G(s)\|_2 = \sqrt{\frac{1}{2\pi}\int_{-\infty}^{\infty} |G(j\omega)|^2 d\omega} = \sqrt{\frac{100^2}{2\pi}\int_{-\infty}^{\infty} \frac{1}{(\omega^2+1)^2} d\omega} = \sqrt{\frac{100^2}{2\pi}\frac{\pi}{2}} = 50.$$

The cumulative integral across frequencies (doubled since we integrate over positive frequencies) is shown in the figure, which asymptotes to the $\mathcal{H}_2$ value. Not only does this network induce a 100-fold maximum amplification of energy-bounded shocks, with standard Gaussian white-noise shocks it induces a 50-fold amplification of the output standard deviation.

Below I provide some general formulas for the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ norms, as well as explicit examples in different settings. But first I provide more discussion on how the $\mathcal{H}_\infty$ norm provides a measure of the system robustness.

3.7 Uncertainty and Robustness

Dynamic networks are subject to a wide range of uncertainties. These may arise from changes or misspecification in the network itself, or in how outputs and inputs are measured

or interact. By treating the input $u(t)$ as an uncertain input, the $\mathcal{H}_\infty$ norm of the system transfer function provides a measure of the robustness of the system to misspecification.

A general and flexible approach to modeling uncertainty is to represent misspecification as a feedback distortion of the effective disturbance entering the model. This means that the input $u(s)$ can consist both of specified shocks $u_0(s)$ and unknown modeling errors or uncertainty $\Delta(s)$ that may feed back on the output. Thus the actual system becomes:

$$y(s) = G(s)(u_0(s) + \Delta(s)y(s)) = G(s)u(s). \tag{6}$$

In Appendix A.2 I show that many forms of uncertainty can be mapped into, or at least bounded by, a representation like (6). This includes uncertainty about the network structure, its parameters or dynamics, uncertainty about the structure or correlations of shocks, or uncertainty in output measurement due to aggregation, delay, or mismeasurement.

A general way to introduce robustness concerns is to consider unstructured uncertainty. I assume that the perturbation $\Delta(s)$ is stable and satisfies a uniform bound of $1/\gamma$ across all frequencies:

$$\|\Delta(s)\|_\infty \equiv \sup_{\omega\in\mathbb{R}} \bar{\sigma}(\Delta(j\omega)) < \frac{1}{\gamma}.$$

This captures the idea that the precise nature of the uncertainty is unknown, but its overall magnitude is limited. A fundamental result from robust control theory is the small gain theorem, which provides a sharp characterization of when a system remains stable in the face of such uncertainty. The theorem, and its proof sketch in Appendix A.3, follow Zhou, Doyle, and Glover (1996).

Theorem 1 *Let $G(s)$ be the nominal transfer function of a linear time-invariant system, and suppose that the uncertainty $\Delta(s) \in \mathcal{H}_\infty$ enters through the feedback interconnection* (6)*. Suppose that $G(s)$ and $\Delta(s)$ are both stable transfer functions. Then the closed-loop system is internally stable for all $\Delta(s)$ satisfying $\|\Delta\|_\infty < 1/\gamma$ if and only if*

$$\|G\|_\infty \leq \gamma.$$

The transfer function $G(s)$ describes how shocks at different frequencies and across different nodes in the network are transmitted to outputs of interest. If the worst-case amplification from input to output is large, say if $\|G(s)\|_\infty$ is close to or exceeds one, then even modest misspecifications or high-frequency distortions in the input can be destabilizing. In contrast, if $G(s)$ uniformly attenuates shocks across all directions and frequencies, the system can tolerate a much broader set of uncertainties without risk of instability. Thus I focus on the $\mathcal{H}_\infty$ norm of systems not only as a way of gauging their performance for arbitrary shocks, but also as a measure of their broader robustness to uncertainty.

The small gain theorem provides a powerful robustness test, but it treats all uncertainty as unstructured. In practice, however, we often have more information about the structure of the uncertainty, as in Hansen and Sargent (2022). In such cases the $\mathcal{H}_\infty$ bound may be overly conservative. Refinements that use structured singular value analysis, also known as μ-analysis as described in Zhou and Doyle (1998), can provide more accurate robustness bounds by incorporating the structure of the uncertainty.

4 Characterization of Performance and Robustness for Network Dynamics

In this section I characterize both the average–case and worst–case amplification in network systems, using the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ norms as described above. Rather than change notation throughout, for a matrix X, let X^H denote the Hermitian transpose when X is complex and the (regular) transpose when X is real.

4.1 General Formulas for the System Norms

4.1.1 $\mathcal{H}_2$ Norm

I make the standard assumption that the input-output map as measured by the transfer function is stable, which ensures that the $\mathcal{H}_2$ norm is finite.[16] Moreover, we can calculate the norm by solving a matrix Lyapunov equation, which effectively computes a steady-state covariance matrix, as shown in Appendix A.4.

Theorem 2 *Suppose M is Hurwitz, and let P be the unique solution of the continuous Lyapunov equation*

$$MP + PM^H + BB^H = 0.$$

Then

$$\|G\|_2^2 = \mathrm{tr}(CPC^H) = \mathrm{tr}(PC^HC). \tag{7}$$

If L, or equivalently M, is diagonalizable, we can write:

$$L = V\Lambda V^{-1}, \qquad M = -L = -V\Lambda V^{-1},$$

[16] When the state matrix M is Hurwitz, as in production networks, then standard controllability and detectabilty conditions suffice. (See Skogestad and Postlethwaite (2005), Chapter 4.) For disagreement dynamics, M is not Hurwitz as it has a zero eigenvalue, but the disagreement projection ensures that the transfer function is stable. Thus we can solve Lyapunov equation on a reduced subspace, or equivalently, work directly with the modal representation which removes the consensus mode.

where $\Lambda = \text{diag}(\lambda_i(L))$. Then define the modal projections

$$B_e = V^{-1}BB^HV^{-H}, \qquad S_e = V^HC^HCV.$$

Here B_e captures how shocks load onto modes, and S_e is how modes map into outputs. Then the Lyapunov equation in modal coordinates has solution $P = VYV^H$, where

$$Y_{ij} = \frac{B_{e,ij}}{\lambda_i + \overline{\lambda_j}}.$$

Thus we have the following result, adapted from Yadykin and Iskakov (2021), whose proof is in Appendix A.5.

Theorem 3 *When L is diagonalizable, the $\mathcal{H}_2$ norm can be represented as:*

$$\|G\|_2^2 = \sum_{i,j} \frac{S_{e,ji}B_{e,ij}}{\lambda_i + \overline{\lambda_j}}. \tag{8}$$

This expression shows explicitly how the norm depends on all of the modes and on the projections of B and C into the basis of the eigenvectors. Non-normality (non-orthogonal eigenvectors) affects B_e and S_e, hence the weights placed on modal interactions. Slow modes, whose eigenvalues have small real parts, contribute disproportionately to the $\mathcal{H}_2$ norm.

4.1.2 The $\mathcal{H}_\infty$ Norm

As above, the $\mathcal{H}_\infty$ norm measures the worst–case amplification of bounded–energy inputs, and can be written

$$\|G\|_\infty = \sup_{\omega\in\mathbb{R}} \bar{\sigma}\big(G(j\omega)\big) = \sup_{\omega\in\mathbb{R}} \bar{\sigma}\big(C(j\omega I - M)^{-1}B\big)\,.$$

The worst–case disturbance, both in direction and frequency, is the one whose image under the resolvent $(j\omega I - M)^{-1}$ is most strongly amplified by B and C. While the $\mathcal{H}_2$ norm reduces to a Lyapunov equation, here we do not have a closed form expression, but rather must take a supremum across frequencies.[17]

If L is diagonalizable, denote the left and right eigenvalue-eigenvector pairs as $Lv_i = \lambda_i v_i$ and $w_i^H L = \lambda_i w_i^H$, where we normalize $w_i^H v_i = \delta_{ij}$, the Kronecker delta. Then via Sylvester's formula for analytic functions we can write the resolvent expansion[18]

$$(sI - M)^{-1} = (sI + L)^{-1} = \sum_i \frac{v_i w_i^H}{s + \lambda_i}.$$

[17] There is a control-theoretic solution where the existence of value for a zero-sum game, which is found by solving a Riccati equation, ensures $\|G\|_\infty < \gamma$. See Zhou, Doyle, and Glover (1996).

[18] See Lancaster and Tismenetsky (1985), Chapter 9.

Consequently we have the spectral expansion of the transfer function:

$$G(s) = C(sI - M)^{-1}B = \sum_i \frac{Cv_i w_i^H B}{s + \lambda_i}. \tag{9}$$

Therefore we can write the $\mathcal{H}_\infty$ norm as:

$$\|G\|_\infty = \sup_{\omega \in \mathbb{R}} \bar{\sigma}\left(\sum_i \frac{Cv_i w_i^H B}{j\omega + \lambda_i} \right). \tag{10}$$

This form makes explicit the contribution of each eigenvalue to the transfer function, with the numerator measuring how strongly that mode affects the input–output map.

I now show how the general formulas for the system norms simplify in some special cases.

4.2 Positive Systems

When it comes to the $\mathcal{H}_\infty$ robustness properties of positive, stable systems like the production network, the zero-frequency behavior matters most, as illustrated in Figure 1. I now show formally that for stable systems with nonnegative A, B, and C, the $\mathcal{H}_\infty$ norm always peaks at frequency zero. The result is adapted from Rantzer (2012), with proof in Appendix A.6.

Theorem 4 *Let the adjacency matrix $A \geq 0$ have row sums $\sum_j A_{i,j} = \beta_i < 1$, and assume $B \geq 0$ and $C \geq 0$. Then*

$$\|G\|_\infty = \sup_{\omega \in \mathbb{R}} \|G(j\omega)\|_2 = \|G(0)\|_2 = \|CL^{-1}B\|_2 = \bar{\sigma}(CL^{-1}B).$$

With this result we can also diagnose the worst case inputs and their impacts on outputs. Let $\bar{v}$ and $\bar{u}$ be the right and left singular vectors corresponding to $\bar{\sigma}(CL^{-1}B)$. Then the worst-case input direction is $\bar{v}$:

$$\sup_\omega \max_{\{u : \|u\|_2 = 1\}} \|G(j\omega)u\|_2 = \|G(0)\bar{v}\|_2.$$

The corresponding most sensitive output pattern is $\bar{u}$:

$$G(0)\bar{v} = \bar{\sigma}(CL^{-1}B)\,\bar{u}.$$

Recall that in Section 3.3 we showed that when $B = C = I$, then $G(0) = L^{-1} = (I - A)^{-1}$. This is the Leontief inverse, the fundamental object of interest in static production network models. Theorem 4 shows that the worst-case amplification of disturbances in the dynamic system is governed precisely by the largest singular value of the Leontief inverse. The corresponding singular vectors identify the most destabilizing direction in input space and the most sensitive pattern in output space.

The average performance of the production network, as measured by the $\mathcal{H}_2$ norm, in general depends on the system dynamics and intermediate frequencies. But the robustness properties are governed by long-run, steady-state responses at frequency zero. The robustness of the production network is therefore determined by the spectral properties of the Leontief inverse. I use this result below for different production network examples.

4.3 Normal Matrices and Full State Analysis

We expressed the system norms in modal form in (8) and (10) for general directed networks. But when the state dynamics are normal ($MM^H = M^H M$), we can simplify these expressions. Recall that undirected networks, as well as simple cycles and rotations, have normal dynamics. I focus here on the case where $B = I$ so that all states or nodes are subject to shocks or inputs. I separately consider the production network with stable (dissipative) dynamics and the information network with disagreement dynamics.

Since L is normal, we can now write the diagonalization with a unitary matrix V:

$$L = V\Lambda V^H.$$

Hence we may choose an orthonormal eigenbasis:

$$v_i = Ve_i, \qquad w_i = v_i,$$

so that $w_i^H v_j = \delta_{ij}$.

4.3.1 Normal Production Networks, $B = C = I$

Here I focus on production networks with positive, stable dynamics and set $C = I$, and so measure the implications for the full output vector. A special case is when L is real and symmetric, as in the case of an undirected production network.

I start with the $\mathcal{H}_2$ formula (8). With normal dynamics and $B = C = I$, the modal input and output weights become

$$B_{e,ij} = w_i^H BB^H w_j = w_i^H w_j = \delta_{ij}, \qquad S_{e,ji} = v_j^H C^H C v_i = v_j^H v_i = \delta_{ij},$$

so $B_e = S_e = I$ in the eigenbasis of L. Substituting into (8) gives

$$\|G\|_2^2 = \sum_{i,j} \frac{\delta_{ij}\,\delta_{ij}}{\lambda_i + \overline{\lambda_j}} = \sum_i \frac{1}{\lambda_i + \overline{\lambda_i}} = \frac{1}{2}\sum_i \frac{1}{\Re\lambda_i}. \tag{11}$$

In particular, when L is real and symmetric (so $\lambda_i \in \mathbb{R}_{++}$), this reduces to

$$\|G\|_2^2 = \frac{1}{2}\sum_{i=1}^N \frac{1}{\lambda_i} = \frac{1}{2}\,\mathrm{tr}(L^{-1}) = \frac{1}{2}\sum_{i=1}^N \frac{1}{1-\lambda_i(A)}.$$

Turning to the $\mathcal{H}_\infty$ case, with $B = C = I$, we have from Theorem 4:

$$\|G\|_\infty = \bar{\sigma}(L^{-1}) = \max_i \frac{1}{|\lambda_i|} = \frac{1}{\min_i |\lambda_i|}. \tag{12}$$

This result can also be derived from the modal representation of the norm in (10). When L is real and symmetric, this simplifies further:

$$\|G\|_\infty = \frac{1}{\min_i \lambda_i} = \frac{1}{1 - \max_i \lambda_i(A)}.$$

For undirected production networks, both the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ norms depend on the inverses of the eigenvalues of the Laplacian, or correspondingly the inverse of the distance between the eigenvalues of the adjacency matrix and one. For the $\mathcal{H}_2$ norm, all of the Laplacian eigenvalues matter, and the norm value is proportional to the sum of their inverses. For the $\mathcal{H}_\infty$ norm, only the smallest Laplacian eigenvalue, or largest adjacency eigenvalue, matters. Clearly as the system approaches instability, with smallest Laplacian eigenvalue approaching zero or largest adjacency eigenvalue approaching unity, both norms explode.

4.3.2 Normal Disagreement Dynamics, $B = I, C = \check{C}$

I now consider the consensus dynamics from Section 2.2, with a row-stochastic and normal adjacency matrix A. Hence the Laplacian has a zero eigenvalue. I continue to assume $B = I$, but now analyze disagreement dynamics with $C = \check{C}$ from (1).

For the modal $\mathcal{H}_2$ formula (8), with normal dynamics the input weights are $B_e = I$ as in the previous section. Now, because $Cv_1 = \check{C}1_N = 0$, we have:

$$S_{e,1i} = v_1^H C^H C v_i = 0, \qquad S_{e,i1} = v_i^H C^H C v_1 = 0,$$

for all i. The matrix $\check{C}$ projects the dynamics onto the orthogonal complement of 1_N, and the eigenvectors $v_2, \ldots, v_N$ form an orthonormal basis of this subspace. So $S_e = \text{diag}(0, 1, \ldots, 1)$ and $S_{e,ii} = 1$ for $i \geq 2$. Substituting these into (8), all the cross-terms drop out, along with the contribution of the first (consensus) mode:

$$\|G\|_2^2 = \sum_{i=2}^{N} \frac{1}{\lambda_i + \overline{\lambda_i}} = \frac{1}{2} \sum_{i=2}^{N} \frac{1}{\Re\lambda_i}. \tag{13}$$

As in the previous section, when L is real and symmetric, this simplifies to:

$$\|G\|_2^2 = \frac{1}{2} \sum_{i=2}^{N} \frac{1}{\lambda_i} = \frac{1}{2} \sum_{i=2}^{N} \frac{1}{1 - \lambda_i(A)}.$$

The disagreement dynamics filter out the neutrally stable consensus mode, but otherwise the $\mathcal{H}_2$ norm is the same as for normal production networks, half the sum of the inverses of the eigenvalues of the Laplacian.

For the $\mathcal{H}_\infty$ norm case, we start with the spectral representation of the transfer function (9). With $v_i = w_i$, $B = I$, and $C = \check{C}$, this becomes:

$$G(s) = \sum_{i=1}^{N} \frac{C v_i v_i^H}{s + \lambda_i} = \sum_{i=2}^{N} \frac{v_i v_i^H}{s + \lambda_i}.$$

Thus we can write

$$G(j\omega) = V \text{diag}\Big(0, \frac{1}{j\omega + \lambda_2}, \ldots, \frac{1}{j\omega + \lambda_N}\Big) V^H.$$

Since V is unitary, the transfer function shares the singular values with the central diagonal matrix. The nonzero singular values are

$$\sigma_i\big(G(j\omega)\big) = \left|\frac{1}{j\omega + \lambda_i}\right|, \qquad i \geq 2.$$

In general, the eigenvalues of the Laplacian are complex, so we can write them as: $\lambda_i = a_i + jb_i$, with $a_i = \Re\lambda_i > 0$. Then note that we have:

$$\sup_{\omega\in\mathbb{R}} \sigma_i\big(G(j\omega)\big) = \sup_{\omega\in\mathbb{R}} \left|\frac{1}{j\omega + \lambda_i}\right| = \sup_{\omega\in\mathbb{R}} \frac{1}{\sqrt{a_i^2 + (\omega + b_i)^2}} = \frac{1}{\sqrt{a_i^2}} = \frac{1}{\Re\lambda_i}. \tag{14}$$

The $\mathcal{H}_\infty$ norm is the largest singular value over disagreement modes:

$$\|G\|_\infty \;=\; \max_{i\geq 2} \frac{1}{\Re\lambda_i} = \frac{1}{\min_{i\geq 2} \Re\lambda_i}. \tag{15}$$

Again, when L is real and symmetric, this reduces to:

$$\|G\|_\infty = \frac{1}{\min_{i\geq 2} \lambda_i} = \frac{1}{\lambda_2} = \frac{1}{1 - \lambda_{N-1}(A)},$$

where we follow the convention of ordering eigenvalues from smallest to largest.

Thus the $\mathcal{H}_\infty$ norm is the inverse of the second smallest eigenvalue of the Laplacian, or second largest eigenvalue of the adjacency matrix. The Laplacian eigenvalue λ_2 is known as the algebraic connectivity of the graph (or network), and it governs the speed of learning in network models.[19] Low algebraic connectivity means slow diffusion across the network and slow learning, hence large amplification of inputs (shocks) to outputs (belief differences).

We have seen that positive systems, like production networks, have their peak response at frequency zero, and hence their worst-case inputs are essentially steady-state perturbations. But because we build in long-run consensus as the steady state of information propagation, disagreement dynamics can give rise to peaks at interior frequencies. In particular, in (14)

[19]See Van Mieghem (2023) for discussion of the algebraic connectivity in graphs and Golub and Jackson (2012) and Golub and Sadler (2016) for its application to the speed of learning.

we saw that for each eigenvalue λ_i the minimizing frequency is $\omega_i^* = -b_i = -\Im\lambda_i$. Thus the overall peak frequency is attained at the negative of the imaginary part of the appropriate eigenvalue. That is, we let $i^* = \arg\min_{i\geq 2} \Re\lambda_i$, and hence $\omega^* = -\Im\lambda_{i^*}$. The worst-case input vector is the right singular vector at the frequency ω^*, which in this case is simply the eigenvector v_{i^*}:

$$G(j\omega^*)v_{i^*} = \frac{1}{j\omega^* + \lambda_{i^*}} v_{i^*} = \frac{1}{a_{i^*}} e^{-j\phi_{i^*}} v_{i^*},$$

where the phase ϕ_{i^*} provides an irrelevant scaling. Hence the worst-case inputs can lead to amplified disagreement for an extended but not permanent duration. But this relies on the Laplacian having complex eigenvalues. When L is real and symmetric, clearly $\Im\lambda_i = 0$ for all i and hence the maximum amplification is at frequency $\omega^* = 0$.

4.4 Production Networks and Aggregate Output

In some situations when analyzing production networks, we care more about aggregate output than the whole cross-sectional output vector. I keep $B = I$ but then instead of taking $C = I$, set $C = 1_N^\top$, so the system output is $y_t = \sum_{i=1}^N x_{i,t}$. The approach here extends to any constant vector, giving a weighted sum of the sectoral outputs. I first show that the $\mathcal{H}_\infty$ norm has some interpretable implications for aggregate output in general production networks. Then I specialize to normal production networks, where both performance measures simplify. In particular, focusing on aggregate output leads to a focus on different features of the network structure compared to analysis of the full cross-sectional output vector.

4.4.1 The $\mathcal{H}_\infty$ Norm in General

Since aggregate output is a scalar, the transfer function becomes a vector-valued function:

$$G(s) = 1_N^\top (sI + L)^{-1}.$$

Then $G(0) = 1_N^\top L^{-1}$ is a $1 \times N$ row vector with nonnegative elements, which we can write as $G(0) = [G_{0,j}]$ where:

$$G_{0,j} = \sum_{i=1}^N \left(L^{-1}\right)_{i,j}.$$

Then the $\mathcal{H}_\infty$ norm is just the Euclidean norm of this vector:

$$\|G\|_\infty = \|G(0)\|_2 = \|1_N^\top L^{-1}\|_2 = \sqrt{\sum_{j=1}^N \left(\sum_{i=1}^N \left(L^{-1}\right)_{i,j}\right)^2}.$$

Since $G(0)$ is a real vector, $\|G(0)\|_2$ is its only singular value, and the right singular vector is proportional to $L^{-\top} 1_N$, the transpose of $G(0)$.

Recall that L^{-1} is the Leontief inverse of the production network. The inner sum in the expression above, $G_{0,j}$, is the total economy-wide output response to a unit shock applied at node j. In static production networks, Domar weights capture the response of output to sectoral productivity changes (see Carvalho and Tahbaz-Salehi (2019)). $G_{0,j}$ plays a similar role here, but now measures the aggregate output response to persistent sectoral shocks. The outer sum in the evaluation of the norm then squares these weights and adds them up, providing a measure of the aggregate worst-case amplification of shocks across nodes. If shocks can hit multiple nodes in the network in correlated ways, the worst possible aggregate outcome occurs when the shock aligns with the vector of per-node aggregate impacts.

4.4.2 Normal Dynamics

When the production network Laplacian L is normal, we can simplify the expressions for both performance measures. With the same diagonalization as in the previous normal case above, $L = V\Lambda V^H$, define the coordinates of the aggregate output vector 1_N in the eigenbasis of L by

$$\alpha = V^H 1_N, \text{ with elements } \alpha_i = v_i^H 1_N.$$

Then, using the $\mathcal{H}_2$ modal formula (8), we again have $B_e = I$ and $B_{e,ij} = \delta_{ij}$, but now:

$$S_{e,ji} = v_j^H C^H C v_i = v_j^H 1_N 1_N^\top v_i = (1_N^\top v_j)^* (1_N^\top v_i) = \overline{\alpha}_j \alpha_i,$$

Then we have:

$$\|G\|_2^2 = \sum_{i,j} \frac{\delta_{ij} |\alpha_i|^2}{\lambda_i + \overline{\lambda_j}} = \sum_i \frac{|\alpha_i|^2}{\lambda_i + \overline{\lambda_i}} = \frac{1}{2} \sum_i \frac{|\alpha_i|^2}{\Re \lambda_i}. \tag{16}$$

Further, when L is symmetric, by the spectral decomposition $L^{-1} = \sum_i \lambda_i^{-1} v_i v_i^\top$, hence

$$\|G\|_2^2 = \sum_i \frac{|\alpha_i|^2}{\lambda_i} = \sum_i \lambda_i^{-1} (1_N^\top v_i)^2 = 1_N^\top L^{-1} 1_N.$$

The expression (16) is similar to (11) for the full output vector. But now the impact of each eigenvalue λ_i is weighted by $|\alpha_i|$, which measures how aligned mode i is with aggregate output. As before, slow-decaying modes with small $\Re \lambda_i$ generate large internal amplification. But now if $|\alpha_i|$ is large, then eigenvector v_i involves many firms responding in the same direction or in a way that has a large influence on aggregate output. So the impact of a mode on the volatility of aggregate output depends both on how it impacts sectoral volatility, and how that sectoral volatility feeds into aggregate volatility.

For the $\mathcal{H}_\infty$ norm, we can use the diagonalization to specialize:

$$G(0) = 1_N^\top L^{-1} = 1_N^\top V \Lambda^{-1} V^H = \alpha^H \Lambda^{-1} V^H.$$

Then taking the norm of this vector leads to the result:

$$\|G\|_\infty^2 = \alpha^H \text{diag}\Big(\frac{1}{|\lambda_i|^2}\Big)\alpha = \sum_{i=1}^{n} \frac{|\alpha_i|^2}{|\lambda_i|^2}. \tag{17}$$

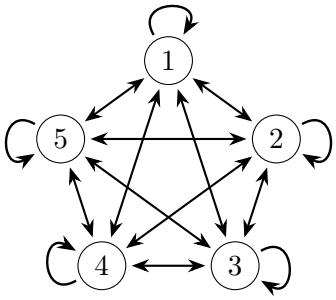
When analyzing aggregate output, the two measures of system performance are more similar than when focused on the full cross-sectional output vector. Here both depend on all of the system modes through the sums of the ratios of $|\alpha_i|$, the mode's contribution to aggregate output, to a feature of the corresponding eigenvalue. For $\mathcal{H}_2$ the real part of the eigenvalue is what matters for overall volatility, while for $\mathcal{H}_\infty$ the magnitude of the eigenvalue is what matters for maximum amplification. When analyzing the full output cross-section with $C = I$ above, the worst-case disturbance is concentrated on the least-stable mode. But when aggregating from the cross-section, every mode matters, depending on how slowly it decays and how much it contributes to aggregate output, even in the worst case.

5 Illustrative Examples

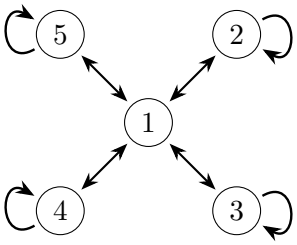
In this section I apply the analysis derived above to several canonical network structures. I illustrate how the average-case ($\mathcal{H}_2$) and worst-case ($\mathcal{H}_\infty$) amplification properties depend on the network topology. For each network, I focus on three applications: cross-sectional output in production networks, aggregate output in the same setting, and disagreement dynamics in information propagation. In each example, I start with a specific row-stochastic adjacency matrix A_0, and in all cases I set $B = I$. For disagreement dynamics, the network is simply $A = A_0$ with Laplacian $L = I - A_0$, and I set $C = \check{C}$.

For production networks, a natural starting point is to assume that A_0 governs the relative allocation of intermediate inputs, where the labor inputs serve as a dampening factor which stabilizes the system. If all sectors have the same intermediate share parameter $\beta_i = \beta \in (0, 1)$, we can write the production network as $A = \beta A_0$ and $L = I - A = I - \beta A_0$. However this simple case leads to some stark results, where performance metrics are independent of the network structure. Thus I also consider a variation with heterogeneity in input shares. For production networks, I consider both cross-sectional output where $C = I$ and aggregate output where $C = 1_N^\top$.

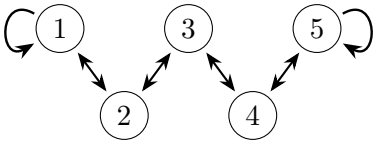
I start with some examples of normal graphs, where the spectrum of the Laplacian can be calculated explicitly. These allow me to illustrate analytically how the performance measures depend on the network properties. I also isolate the worst-case frequencies and worst-case inputs, which show where the networks are most fragile. Even in the normal case, I illustrate how directed graphs can shift worst-case amplification toward interior frequencies, for example amplifying disagreement. Then I turn to non-normal cases, showing how the Laplacian spectrum no longer characterizes amplification or network fragility.



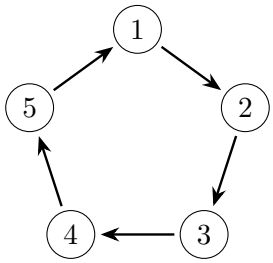
(a) Complete: self-loop at each node.



(b) Star: hub at 1, self-loops at leaves only.



(c) Path: symmetric with endpoint loops only.



(d) Directed cycle: asymmetric, no self-loops.

Figure 3: Example normal network configurations with $N = 5$. A link $j \to i$ corresponds to $A_{ij} > 0$. Self-loops are drawn when $A_{ii} > 0$.

5.1 Examples with Normal Dynamics

With normal dynamics we can reduce the calculation of the performance norms to components of the Laplacian spectrum. The following examples use characterizations in the literature of the spectrum of some classical graphs. For normal production networks, I apply these results in (11) and (12) for cross-sectional output and (16) and (17) for aggregate output. With disagreement dynamics, I use the formulas (13) and (15).

The particular examples that I study in this section are: complete networks, star networks, paths, and directed cycles. These configurations are illustrated for $N = 5$ in Figure 3. Recall the convention that a link $j \to i$ corresponds to $A_{ij} > 0$. Self-loops are drawn when $A_{ii} > 0$, which is necessary in some cases to ensure that the adjacency matrices are row-stochastic and/or symmetric. The directed cycle is normal but asymmetric. With a complete analytic characterization of the spectrum of the Laplacian, I have explicit results for the performance metrics. The examples differ in their degree of connectedness, centrality, and redundancy of links and so are useful to illustrate how the network topology affects the propagation of shocks, and can lead to bottlenecks and amplification.

5.1.1 The Role of Labor in Production Networks

The simplest normal production networks are symmetric, with a homogeneous intermediate input share β. Since the baseline matrix A_0 has a unit eigenvalue, the largest eigenvalue of $A = \beta A_0$ is β. Therefore, in all cases regardless of the details of the network structure, the $\mathcal{H}_\infty$ norm in (12) is simply

$$\|G\|_\infty = \frac{1}{1-\beta}.$$

Further, $v_1 = \frac{1}{\sqrt{N}} 1_N$ is the eigenvector associated with the Laplacian eigenvalue $1-\beta$, and since L is normal, the other eigenvalues are orthogonal to v_1. Hence we have $\alpha_1 = v_1^\top 1_N = \sqrt{N}$ while $\alpha_i = v_i^\top 1_N = 0$ for $i = 2, \ldots, N$. For aggregate output, (16) and (17) reduce to:

$$\|G\|_2^2 = \frac{N}{2(1-\beta)}, \qquad \|G\|_\infty = \frac{\sqrt{N}}{1-\beta}.$$

Thus these performance metrics only depend on the size of the network and the intermediate input share β, but not on any other aspect of the network structure.

For production networks, labor (or other non-networked inputs) gives a type of leakage from the network that dampens the impact of shocks. For symmetric networks with identical labor shares, the dampening is uniform. Thus for $\mathcal{H}_\infty$, the worst-case amplification occurs with a uniform shock to all sectors, giving simple expressions for the cross-sectional and aggregate amplification. For the $\mathcal{H}_2$ norm, all modes of the network matter and affect cross-sectional output volatility. But for aggregate output, the idiosyncratic effects of the individual sectors largely cancel. The network propagates only the common component of shocks, which is uniformly damped by labor. Thus a sparse production network can generate substantial sectoral dispersion and slow reallocation dynamics, while still leaving aggregate volatility comparatively muted.

For production networks, non-normal dynamics arise generically. Even when input–output relations are otherwise symmetric, heterogeneity in input shares across sectors typically renders the resulting dynamics non-normal. However, there is a convenient specification of input share heterogeneity which preserves the normality of the dynamics. I start with a reference value β, then define a perturbation vector $\Gamma = \mathrm{diag}(\beta_1 - \beta, \ldots, \beta_N - \beta)$, and set the adjacency matrix to:

$$A = \beta A_0 + \Gamma.$$

This specification introduces heterogeneous input shares, β_i, by altering the diagonal of the adjacency, meaning that sectors with lower input shares use less of their own goods as inputs. This specification preserves normality, since if A_0 is symmetric then A is as well.

The Laplacian then becomes:

$$\begin{aligned} L &= I - A = I - \beta A_0 - \Gamma = D - \beta A_0, \text{ where:} \\ D &= \text{diag}(d_1, \ldots, d_N), \quad d_i = 1 - \beta_i + \beta. \end{aligned}$$

In the following I show how input share heterogeneity affects network dynamics by changing the effective weights of different modes, and hence the local speed of adjustment.

5.1.2 Complete Network

I start with the symmetric fully-connected complete network, which illustrates the extreme of connectedness and redundancy, as each node is connected with every other node. This is a representation of thick interactions, with multiple suppliers and buyers in the production context, and diffuse propagation of information in the learning context. Here the baseline adjacency matrix is

$$A_0 = \frac{1}{N} 1_N 1_N^\top.$$

The baseline Laplacian $L_0 = I - A_0$ then has eigenvalues $\lambda_1 = 0$, $\lambda_2 = \cdots = \lambda_N = 1$.[20]

5.1.2.1 Disagreement Dynamics

Thus for the disagreement dynamics, (13) and (15) yield:

$$\|G\|_2^2 = \frac{1}{2} \sum_{i=2}^{N} \frac{1}{\lambda_i} = \frac{N-1}{2}, \quad \|G\|_\infty = \frac{1}{\lambda_2} = 1.$$

Thus the average performance of the network scales with the square root of its size, as adding sectors adds more sources of idiosyncratic shocks. With more nodes, there are more distinct ways for shocks to create temporary dispersion across the network. By contrast, the maximum amplification remains at one regardless of the size of the network. Since the Laplacian eigenvalues are all real, the worst case is at frequency zero. Moreover, any unit vector orthogonal to 1_N is an eigenvector for λ_2 and one may choose, for instance,

$$\bar{v} = \frac{1}{\sqrt{2}} \big(1, -1, 0, \ldots, 0\big)^\top$$

as the worst-case disturbance pattern. Every node in the network is equally well connected. So the most damaging perturbation involves shocking one node up and another down by equal and opposite amounts, which leads to the maximal disagreement. But because every node or sector is tightly and symmetrically linked, that disagreement is quickly averaged out and not amplified.

[20] Van Mieghem (2023) provides related results under a slightly different normalization of the graph.

5.1.2.2 Production Networks: Cross-Section

For production networks, we start with cross-sectional output $C = I$ and homogeneous intermediate shares β. Using $L = I - \beta A_0 = (1-\beta)I + \beta L_0$, the eigenvalues of Laplacian are then $\lambda_1 = 1 - \beta$, $\lambda_2 = \cdots = \lambda_N = 1$. From above, we know that the $\mathcal{H}_\infty$ norm is $1/(1-\beta)$ and the worst-case input is a uniform perturbation across sectors at frequency zero. For average performance we use (11):

$$\|G\|_2^2 = \frac{1}{2}\sum_{i=1}^{N}\frac{1}{\lambda_i} = \frac{1}{2}\left(\frac{1}{1-\beta} + N - 1\right).$$

Thus cross-sectional volatility depends on the common component $1/(1-\beta)$ as in the worst-case amplification, but also on the idiosyncratic components which are orthogonal to it. As the size of the network grows, these idiosyncratic components contribute proportionally to overall volatility.

5.1.2.3 Production Networks: Cross-Section, Heterogeneous Labor Shares

Introducing heterogeneous labor shares as above, the Laplacian now becomes:

$$L = I - \beta A_0 - \Gamma = D - \beta A_0 = D - \frac{\beta}{N} 1_N 1_N^\top.$$

Recall that $d_i = 1 - \beta_i + \beta$, and define the following:

$$s_1 = \sum_{i=1}^{N}\frac{1}{d_i}, \qquad s_2 = \sum_{i=1}^{N}\frac{1}{d_i^2}, \qquad \eta = 1 - \frac{\beta}{N}s_1.$$

The complete network with heterogeneous labor shares can be decomposed into a uniform mixing provided by the network, along with different degrees of persistence in each sector. A small intermediate share β_i means a large d_i and therefore more rapid adjustment in that sector. Mathematically, L is composed of a diagonal matrix D plus a rank-one perturbation, so its inverse follows directly from the Sherman–Morrison-Woodbury formula[21]

$$L^{-1} = D^{-1} + \frac{\beta}{\eta N} D^{-1} 1_N 1_N^\top D^{-1}.$$

Therefore

$$\|G\|_2^2 = \frac{1}{2}\operatorname{tr}(L^{-1}) = \frac{1}{2}\left(s_1 + \frac{\beta}{\eta N} s_2\right).$$

The first term in s_1 aggregates sector-level persistence induced by heterogeneous labor shares, capturing how long shocks survive locally before being dissipated. The second term captures the network feedbacks. β/N reflects the strength of uniform mixing in the complete network,

[21] See Horn and Johnson (2013), p. 19.

while η measures the distance to the stability boundary of the aggregate mode. The s_2 term puts extra weight on the most persistent sectors. A sector with high persistence accumulates its own volatility and disproportionately influences aggregate dynamics when sectors interact.

For worst-case analysis, we must find the minimal eigenvalue of L, which is no longer directly related to that of L_0. For $\lambda \notin \{d_1, \ldots, d_n\}$, we write the characteristic equation using the matrix determinant lemma:

$$\det(L - \lambda I) = \det\Big(D - \lambda I - \frac{\beta}{N} 1_N 1_N^\top\Big) = \det(D - \lambda I)\, \det\Big(I - \frac{\beta}{N}(D - \lambda I)^{-1} 1_N 1_N^\top\Big).$$

Since $1_N 1_N^\top$ is rank one,

$$\det\Big(I - \frac{\beta}{N}(D - \lambda I)^{-1} 1_N 1_N^\top\Big) = 1 - \frac{\beta}{N} 1_N^\top (D - \lambda I)^{-1} 1_N = 1 - \frac{\beta}{N} \sum_{i=1}^{N} \frac{1}{d_i - \lambda}.$$

Hence the eigenvalues are given by the roots of

$$1 - \frac{\beta}{N} \sum_{i=1}^{N} \frac{1}{d_i - \lambda} = 0, \tag{18}$$

together with any eigenvalues $\lambda_i = d_i$ that may arise from eigenvectors orthogonal to 1_N. In particular, the smallest eigenvalue $\underline{\lambda} = \min_i \lambda_i$ is the smallest root of the secular equation (18). This equation determines the effective rate at which a common disturbance is dissipated once heterogeneity in labor shares is taken into account. Thus the $\mathcal{H}_\infty$ norm measures the maximum steady-state amplification of a persistent shock that is aligned with this slow aggregate mode. It captures how heterogeneous sectoral adjustment rates interact with uniform linkages to slow down the dissipation of shocks. When some sectors have high labor shares (low β_i, high d_i), the aggregate mode becomes more persistent, increasing worst-case amplification even though the underlying network remains symmetric.

Moreover, the corresponding eigenvector has an explicit componentwise form. From $(L - \lambda I)v = 0$ we have

$$(D - \lambda I)v = \frac{\beta}{N} 1_N (1_N^\top v).$$

so $\overline{v}_i = \frac{c}{d_i - \underline{\lambda}}$ for some constant $c \neq 0$. Sectors with low labor shares (low d_i) receive disproportionately large weight in the worst-case disturbance. The most damaging disturbance is not uniform across sectors once labor shares differ. Instead, it concentrates on those sectors where shocks dissipate slowly. The network then redistributes these persistent local disturbances evenly across the economy through the complete mixing structure.

These results are illustrated in Figure 4, which considers a 20-node complete production network with homogeneity and heterogeneity of labor shares. In the left and center panels the homogeneous case has $\beta = 0.6$ while under heterogeneity the first 10 sectors have $\beta_i = 0.85$

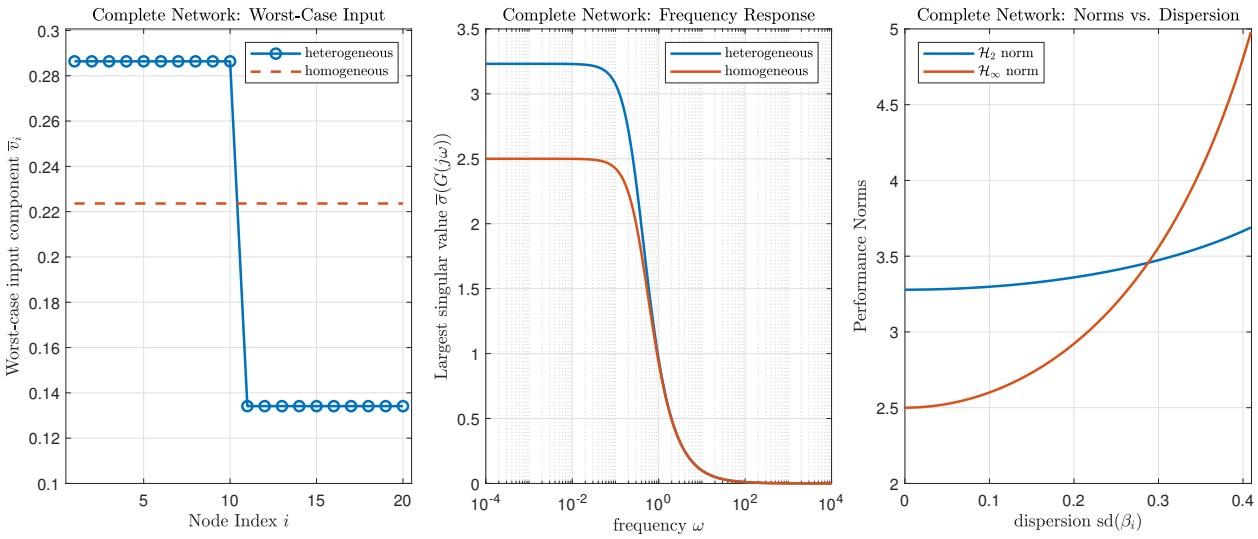


Figure 4: Left panel: Worst-case input vector $\overline{v}$ for a 20-node complete production network with heterogeneous and homogeneous labor shares. Center panel: largest singular values for the complete networks under heterogeneity and homogeneity. The homogeneous case has $\beta = 0.6$ while under heterogeneity 10 sectors have $\beta_i = 0.85$ and 10 have $\beta_i = 0.35$. Right panel: $\mathcal{H}_2$ and $\mathcal{H}_\infty$ performance norms for compete networks with the two-type heterogeneity with varying dispersion.

and the last 10 have $\beta_i = 0.85$, and hence the same mean. The left panel shows the elements of the worst-case input vector $\overline{v}$, while the center panel shows the largest singular value $\overline{\sigma}(G(j\omega))$. In the homogeneous case, the worst-case input is a constant vector and the $\mathcal{H}_\infty$ norm is $1/(1-\beta) = 2.5$. In this network, all connections are complete and symmetric across sectors, so the inputs are targeted at the sectors, with no spill-overs across sectors. Under heterogeneity, the worst-case input is constant across each type of sector (sharing a common labor share parameter), but now concentrates mass in the bottleneck sectors with larger β_i.[22] The center panel shows that heterogeneity increases both the maximum amplification of shocks, as the peak of the singular values increases, as well as overall volatility, as the area under the curve increases.[23] The right panel shows this result more broadly, plotting the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ norms for complete networks with this two-type form of heterogeneity, where we fix the mean input share at 0.6 and increase the dispersion of input shares. More sectoral dispersion increases both overall volatility and especially worst-case amplification.

[22]Since the nodes are symmetric, the labeling and ordering are arbitrary. The apparent "regional" concentration is an artifact of ordering the high β_i sectors first. The vector could equally well be represented as a sawtooth if we alternated high-β_i and low-β_i sectors.

[23]Numerically, when we introduce heterogeneity the $\mathcal{H}_\infty$ norm increases from 2.5 to 3.23 and the $\mathcal{H}_2$ norm increases from 3.28 to 3.42.

5.1.2.4 Production Networks: Aggregate Output

With homogeneous intermediate shares $\beta_i = \beta$, both performance measures collapse to values independent of the network structure. But that is no longer the case with heterogeneous labor shares. Using $L^{-1}1_N = \frac{1}{\eta}D^{-1}1_N$, we have from the calculations above:

$$1_N^\top L^{-1}1_N = 1_N^\top D^{-1}1_N + \frac{\beta}{\eta N}(1_N^\top D^{-1}1_N)^2 = s_1 + \frac{\beta}{\eta N}\, s_1^2.$$

Thus using the result (16) above:

$$\|G\|_2^2 = \frac{1}{2}\, 1_N^\top L^{-1}1_N = \frac{1}{2}\left(s_1 + \frac{\beta}{\eta N}\, s_1^2\right).$$

This is similar to the cross-sectional $\mathcal{H}_2$ norm above, but with the second component having s_1^2 instead of s_2. This reflects the fact that here we aggregate before computing volatility.

Similarly, from (17) above we have:

$$\|G\|_\infty = \|L^{-1}1_N\|_2 = \frac{1}{\eta}\|D^{-1}1_N\|_2 = \frac{\sqrt{s_2}}{\eta}.$$

This expression is quite similar to what we had with uniform labor shares, $\sqrt{N}/(1-\beta)$. Now with heterogeneity, η rather than $1-\beta$ measures the effective stability margin, incorporating the different sectoral adjustment speeds. Similarly, $\sqrt{s_2}$ measures the size of the sectoral persistence weights, which collapses to $\sqrt{N}$ in the homogeneous case. Moreover the worst-case input now has the explicit expression:

$$\overline{v} \propto L^{-1}1_N \propto D^{-1}1_N, \quad \text{so:} \;\; \overline{v}_i \propto \frac{1}{d_i}.$$

The most damaging disturbance for aggregate output concentrates on slow-adjusting, low labor share (low d_i) sectors. Since in this example sectoral connections are all complete and symmetric, aggregate robustness is governed by sectoral bottlenecks.

5.1.3 Star Network

While all nodes were very similar in the complete network, now we look at a symmetric star network with a prominent central node. We consider $N \geq 3$ and let node 1 be the center or hub, and nodes $2, \ldots, N$ be the leaves or spokes, each connected only to the hub. Since leaves do not interact directly but only via the hub, this is a representation of an intermediated market with the hub playing a central role in transactions. I show how shocks can then lead to bottlenecks, delayed adjustment, and propagation of disagreement.

The symmetric row-stochastic star network can be written as the block matrix

$$A_0 = \begin{pmatrix} 0 & (1-a)1_{N-1}^\top \\ (1-a)1_{N-1} & aI_{N-1} \end{pmatrix},$$

where I define $a = 1 - \frac{1}{N-1}$. In Appendix B.1, I show how to compute he spectrum of this graph which yields:

$$\lambda_1 = 0,\ \lambda_2 = \ldots = \lambda_{N-1} = \frac{1}{N-1},\ \lambda_N = 1 + \frac{1}{N-1}.$$

5.1.3.1 Disagreement Dynamics

Thus for the disagreement dynamics, (13) and (15) yield:

$$\|G\|_2^2 = \frac{1}{2}\sum_{i=2}^{N}\frac{1}{\lambda_i} = \frac{1}{2}\left((N-2)(N-1) + \frac{N-1}{N}\right)\ ,\quad \|G\|_\infty = \frac{1}{\lambda_2} = N-1.$$

Thus both the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ norms scale with N, suggesting much more propagation and amplification of disagreement than in the complete network. As discussed above, since the algebraic connectivity λ_2 has multiplicity $N-2$, any vector $x = (0, x_2, \ldots, x_N)^\top$ with $\sum_{j=2}^{N} x_j = 0$ is an eigenvector associated with it. In particular, any shock vector which loads positively on one leaf and negatively on another is a worst-case input, such as:

$$\overline{v} = \frac{1}{\sqrt{2}}(0, 1, -1, 0, \ldots, 0)^\top.$$

But any input which induces potentially varying differences among the leaves that net to zero, while keeping the hub unaffected, is a worst-case input.

Since the hub of the star is the most connected node, one might conjecture that shocks to it would propagate the most. Instead, volatility comes from shocking the leaves in opposite directions. The shocks generate disagreement that can only be dissipated by passing through the hub, but each leaf only "listens" to the hub with weight $1/(N-1)$. That creates a hub bottleneck, where leaf-specific disagreement is mediated only through the center, so the force toward consensus in these directions is weak. The $\mathcal{H}_\infty$ norm captures precisely this slow adjustment process, while the $\mathcal{H}_2$ norm sums across all of these disagreement directions. There are $N-2$ leaf-difference modes, each of which have the same slow rate of decay.

5.1.3.2 Production Networks: Cross-Section

For production networks, we again start with homogeneous labor-shares $\beta_i = \beta$, where the Laplacian becomes $L = (1-\beta)I + \beta L_0$, and hence the eigenvalues are

$$\lambda_1 = 1-\beta, \qquad \lambda_2 = \ldots = \lambda_{N-1} = (1-\beta) + \frac{\beta}{N-1},\ \lambda_N = (1-\beta) + \frac{\beta N}{N-1}.$$

The $\mathcal{H}_\infty$ norm is still $1/(1-\beta)$, but for $\mathcal{H}_2$ we now have:

$$\|G\|_2^2 = \frac{1}{2}\left(\frac{1}{1-\beta} + \frac{N-2}{(1-\beta)+\beta/(N-1)} + \frac{1}{(1-\beta)+\beta N/(N-1)}\right).$$

Note that as the number of sectors N gets large this is approximately:

$$\|G\|_2^2 \approx \frac{1}{2}\left(\frac{1}{1-\beta}+\frac{N-2}{1-\beta}+1\right)+o(1)=\frac{N-1}{2(1-\beta)}+O(1)$$

The $\mathcal{H}_2$ norm aggregates variance across modes, and in the star network there are $N-2$ leaf-difference directions with slow damping of shocks. As in the complete network, cross-sectional volatility scales with N. But in the complete network each new node adds about $1/2$ to overall volatility, while here an additional leaf adds approximately $1/(2(1-\beta))$ due to its slow damping. In a complete network, idiosyncratic shocks are rapidly averaged out, while in a star the many leaf-difference disturbances dissipate only through the hub.

5.1.3.3 Production Networks: Cross-Section, Heterogeneous Labor Shares

We now introduce heterogeneous labor shares by allowing the center to have parameter $\beta_1 = \beta_c$ while all leaves share a common parameter $\beta_i = \beta$ for $i \geq 2$. As above, the Laplacian then takes the form $L = D - \beta A_0$ where now $D = \text{diag}(1-\beta_c+\beta, 1, \ldots, 1)$.

In Appendix B.1, I show how to calculate the spectrum in this case. As in the disagreement case, we can define the leaf-difference modes, but now they reflect the leakage from labor. We now have that $1-\beta a$ is an eigenvalue of L with multiplicity $N-2$. The remaining two eigenvalues, given explicitly in the appendix, can be denoted λ_+ and λ_-.

This form of heterogeneity keeps the star's spectral structure, but shifts the spectrum of the Laplacian, which in turn affects the system norms. In particular, we now have:

$$\|G\|_2^2 = \frac{1}{2}\left(\frac{N-2}{1-\beta a}+\frac{1}{\lambda_+}+\frac{1}{\lambda_-}\right).$$

This has much of the same structure as before, with $N-2$ leaf difference modes. However now the minimal eigenvalue now shifts:

$$\|G\|_\infty = \frac{1}{\min_i \lambda_i} = \frac{1}{\lambda_-}.$$

Moreover the worst-case puts positive weight on the hub and uniform weight on each of the leaves. When $\beta_c > \beta$, then the hub is slower adjusting and the worst-case shock puts more weight on it.

5.1.3.4 Production Networks: Aggregate Output

With homogeneous labor shares aggregate amplification is pinned down by the $1-\beta$ independent of network structure. We therefore focus on the two-type case as in the previous section. Now with aggregate output, the leaf-difference modes play no role. The leaf-difference eigenvectors satisfy $Cx = 1_N^\top x = 0$ and hence they drop out for aggregate output, so only the

modes from the 2×2 hub-leaf block matters. In Appendix B.1, I show how to represent this block via 2×2 matrix M_β. Then the aggregate $\mathcal{H}_2$ norm depends only on this block:

$$\|G\|_2^2 = \frac{1}{2}\ 1_N^\top L^{-1} 1_N = \frac{1}{2}\, b^\top M_\beta^{-1} b,$$

where $b = (1, \sqrt{N-1})^\top$. While these expressions look similar, they mask the simplicity: L is $N \times N$, while M_β is 2×2.[24] Heterogeneity matters for aggregate output in a way that is dependent on the network structure, by changing the weights on the hub-leaf block.

For worst-case amplification, recall that L^{-1} is the Leontief inverse, with multiplier vector $G_0 = L^{-1} 1_N$. By symmetry, we can write $G_0 = (g_1, g_\ell, \ldots, g_\ell)^\top$ with a common value across leaves given explicitly in the appendix. Hence we can write the $\mathcal{H}_\infty$ norm as:

$$\|G\|_\infty = \|L^{-1} 1_N\|_2 = \|G_0\|_2 = \sqrt{g_1^2 + (N-1) g_\ell^2},$$

and the worst-case input is proportional to G_0. In this two-type star there are only two distinct multipliers, one for the hub (g_1) and one common to all leaves (g_ℓ). Heterogeneity matters because it breaks the uniform-shock benchmark and the worst-case aggregate disturbance tilts toward the sector type with the larger multiplier. When the hub is relatively fragile (high β_c, hence low labor share), g_1 rises, so hub shocks are disproportionately important for aggregate output. When the hub is relatively dissipative (low β_c), the leaves carry more of the aggregate multiplier, so the worst-case disturbance loads more heavily on the leaf block. Heterogeneity determines whether aggregate volatility is hub-driven or leaf-driven.

5.1.4 Path Graph

The star network incorporated a form of local interaction, as each leaf only interacted with the hub. But the hub played a special role in propagating information and dissipating shocks. I now turn to a path graph, which is an example where all nodes only interact with their neighbors in a longer production chain. This is an example of thin markets and propagation with location effects, as the impact of a bottleneck depends not just on its characteristics but where it sits in the chain.

Let A_0 be the symmetric row-stochastic path graph, or a reflecting random walk on

[24] Given the simplicity, we can also work out the expression for the $\mathcal{H}_2$ explicitly, but its value doesn't yield much additional insight directly

$\{1, \ldots, N\}$:

$$A_0 \;=\; \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 & \cdots & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 & \cdots & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & \cdots & 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

We can fully characterize the spectrum of the baseline Laplacian $L_0 = I - A_0$.[25] The eigenvalues are:

$$\lambda_i = 1 - \cos\Big(\frac{(i-1)\pi}{N}\Big), \qquad i = 1, \ldots, N,$$

so $\lambda_1 = 0$ and the algebraic connectivity satisfies:

$$\lambda_2 = 1 - \cos\left(\frac{\pi}{N}\right) \sim \frac{\pi^2}{2N^2} \text{ as } N \to \infty.$$

Thus the path network has very low connectivity for large N, scaling like N^{-2}, in comparison to the star network which scaled like N^{-1} and the complete network which was constant.

5.1.4.1 Disagreement Dynamics

Thus for the disagreement dynamics, (13) and (15) yield:

$$\|G\|_2^2 = \frac{1}{2}\sum_{i=2}^{N}\frac{1}{\lambda_i} = \frac{1}{2}\sum_{m=1}^{N-1}\frac{1}{1-\cos\left(\frac{m\pi}{N}\right)}, \quad \|G\|_\infty = \frac{1}{\lambda_2} = \frac{1}{1-\cos\left(\frac{\pi}{N}\right)}.$$

A convenient representation for the worst-case direction is

$$\overline{v}_i = v_{2,i} \;\propto\; \cos\Big(\Big(i - \frac{1}{2}\Big)\frac{\pi}{N}\Big),\; i = 1, \ldots, N.$$

The path has only a local interaction structure. Disagreement cannot be canceled quickly by averaging across many neighbors. Instead, it must diffuse one step at a time along the line. Relative to the complete network, the path has many slow, long-wavelength disagreement modes, so the $\mathcal{H}_2$ disagreement energy grows like N^2 rather than N. Compared to the star, it scales at the same rate but is smaller because persistence is spread across the network, rather than being concentrated on $N-2$ block of leaf-differences.

The worst-case input vector $\overline{v}$ is illustrated in Figure 5, where the component for each node i is plotted at $(i - 1/2)/N$. For $N = 3$, we have $\overline{v} = \frac{1}{\sqrt{2}}(1, 0, -1)^\top$ so the worst-case

[25] The cosine spectrum and eigenfunctions for the double-loop (reflecting) path appears, for example, in Proposition 3.7 of Boyd, Diaconis, Parrilo, and Xiao (2005). For Laplacian spectra of paths under other normalizations, see Van Mieghem (2023) or Brouwer and Haemers (2012).

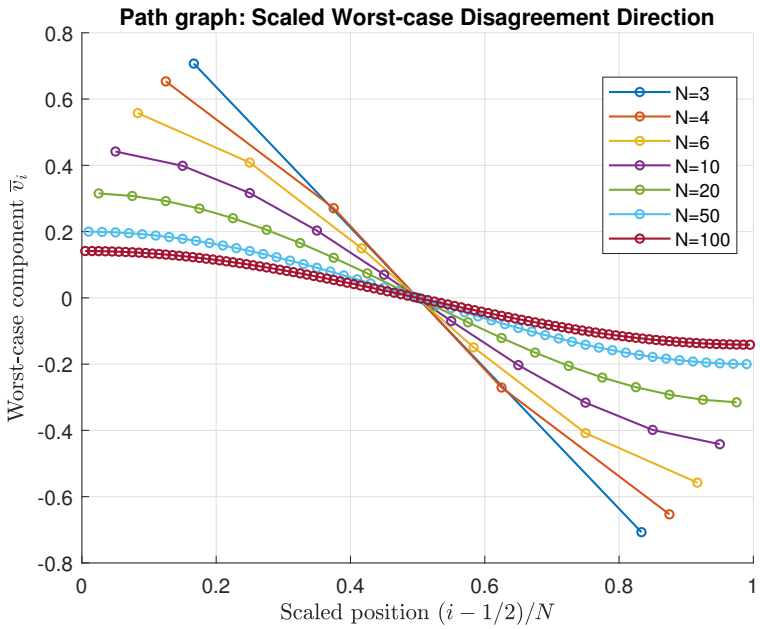


Figure 5: Worst-case input vector $\overline{v}$ for disagreement dynamics on the path graph.

shock pushes one end of the chain up and the other down. As N grows, it is the smooth analogue of such a two-half split: positive on one side of the chain and negative on the other, with sign change near the midpoint. It captures the slowest-mixing imbalance, since undoing the shock requires transfers across the entire chain, diffusing through every link.

5.1.4.2 Production Networks: Cross-Section

With homogeneous labor shares, the Laplacian becomes $L = (1-\beta)I + \beta L_0$ and the eigenvalues are a convex combination of 1 and the eigenvalues of L_0 :

$$\lambda_i = (1-\beta) + \beta\Big(1 - \cos\Big(\frac{(i-1)\pi}{N}\Big)\Big), \qquad i = 1, \ldots, N.$$

As always, the $\mathcal{H}_\infty$ norm collapses to $1/(1-\beta)$, while the $\mathcal{H}_2$ still aggregates the full spectrum:

$$\|G\|_2^2 = \frac{1}{2}\sum_{i=1}^{N} \frac{1}{(1-\beta) + \beta\big(1 - \cos(\frac{(i-1)\pi}{N})\big)}.$$

Even though worst-case amplification is pinned down by the "leakage" $(1-\beta)$ from labor, the path creates many slow spatial modes that contribute disproportionately to $\mathcal{H}_2$. Local shocks propagate gradually through neighboring links, increasing the overall cross-sectional volatility even though the worst case is pinned down by the labor share.

With heterogeneous labor shares, local bottlenecks can develop around the slowest adjusting modes with the lowest labor shares. Since the diffusion here is local, overall volatility depends not only on the distribution of labor shares but also on where they sit along the

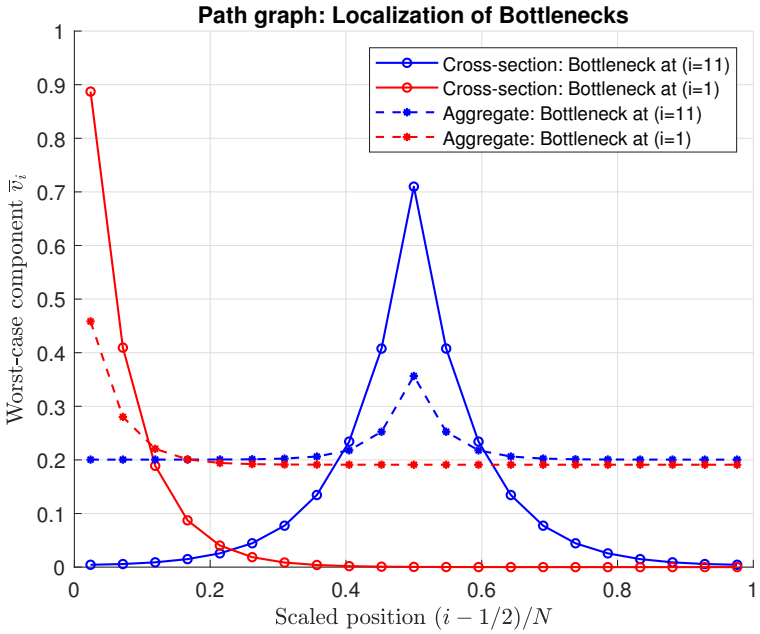


Figure 6: Worst-case input vector $\overline{v}$ for two production networks which differ only in the location of the sector with the smallest labor share. Solid lines are for cross-sectional output, while dashed lines are for aggregate output.

path. Specifically, as above the Laplacian becomes $L = D - \beta A_0$ with $D = \text{diag}(1 - \beta_i + \beta)$, which remains symmetric. While the spectrum is no longer a simple transform of L_0, it has some simple properties given its tridagonal structure. The smallest eigenvalue λ_1 satisfies: $L\overline{v} = \lambda_1 \overline{v}$, which for an interior mode means:

$$-\frac{\beta}{2}\overline{v}_{i-1} + d_i \overline{v}_i - \frac{\beta}{2}\overline{v}_{i+1} = \lambda_1 \overline{v}_i \quad \Rightarrow (d_i - \lambda_1)\overline{v}_i = \frac{\beta}{2}\left(\overline{v}_{i-1} + \overline{v}_{i+1}\right).$$

So when d_i is small (relative to λ_1), then the worst-case input $\overline{v}_i$ tends to be large.

Sectors with smaller than average labor shares, which adjust more slowly, become bottlenecks that contribute more to volatility and thus receive larger weight in the worst-case input. Because interactions are local on a path, volatility depends on the location of the bottlenecks, not only on how many there are. This is illustrated in Figure 6, which shows the components of the worst-case input $\overline{v}$ for two production networks that differ only in the location of the bottleneck sector with the smallest labor share, either in the middle or at one end of the chain.[26] The solid lines in the figure show the cross-sectional analysis here, while the dashed lines are for aggregate output, which is studied in the next section.

No matter where it is located, the bottleneck sector receives the highest weight in the worst-case input. Nearby sectors, which are homogeneous and are only distinguished by

[26]Specifically, the figure uses $N = 21$ sectors, where 20 of them have $\beta_i = 0.6$, but one has $\beta_i = 0.95$.

proximity to the bottleneck, also receive weight that diminishes in distance from the bottleneck sector. When the bottleneck is at the end, shocks can only move up the chain. When the bottleneck sector is in the center of the chain, shocks propagate both up and down the line, which allows the impact of a shock to be spread across more dissipative sectors, reducing both overall volatility and worst-case amplification.[27]

5.1.4.3 Production Networks: Aggregate Output

Again, we focus on the heterogeneous labor share case here. Recall that we have defined the multiplier vector as $G_0 = L^{-1}1_N$, with elements g_i, so that $\|G\|_\infty = \|G_0\|_2$ and $\|G\|_2^2 = \frac{1}{2}\, 1_N^\top G_0$. As in the previous section, the tri-diagonal form of the network facilitates characterization. Here $LG_0 = 1_N$ becomes, for interior nodes:

$$-\frac{\beta}{2}g_{i-1} + d_i\, g_i - \frac{\beta}{2}g_{i+1} = 1 \quad \Rightarrow d_i\, g_i = 1 + \frac{\beta}{2}(g_{i-1} + g_{i+1}).$$

This recursion makes the local spill-over mechanism transparent. A node with low d_i (low labor share, slow local dissipation) carries a larger local multiplier g_i. Because interaction is local, this effect propagates to neighboring multipliers $g_{i\pm1}$ and then decays with distance along the path.

As in the cross-sectional case, aggregate amplification depends not only on the distribution of labor shares but where these sectors lie on the chain. The dashed lines in Figure 6 show the worst-case input for the same example as above, but now focusing on aggregate output. Recall that under homogeneity, the worst-case input is a uniform shock to all sectors. Heterogeneity now tilts the worst-case away from uniformity, and loads more heavily on the bottleneck sector and its neighbors. But since we focus now on the common aggregate component, the shocks are much less concentrated than in the cross-sectional case.

5.1.5 Directed Cycle

We now consider a directed N-cycle, which is the first example where the dynamics are normal but not symmetric. Economically, the directed cycle represents thin input and output markets, where a sector buys intermediates from a single upstream supplier and sells to a single downstream customer. This is the simplest directed graph whose Laplacian has complex eigenvalues, and thus the simplest example in which disagreement dynamics can achieve their worst-case amplification at a strictly positive interior frequency. These phenomena cannot occur in undirected graphs.

The row-stochastic adjacency matrix A_0 shifts mass one step forward around the circle:

$$(A_0)_{i,i-1} = 1 \text{ for } i = 2, \ldots, N, \qquad (A_0)_{1N} = 1,$$

[27] Numerically, with the bottleneck at the end, the $\mathcal{H}_\infty$ norm is approximately 4.73 and $\mathcal{H}_2$ is 3.86. With the bottleneck in the middle, the $\mathcal{H}_\infty$ norm is approximately 3.27 and $\mathcal{H}_2$ is 3.77.

and $(A_0)_{ij} = 0$ otherwise. Thus A_0 is a circulant permutation matrix, hence unitary and therefore normal. Gray (2006) characterizes the spectrum of A_0 in terms of the discrete Fourier transform.[28] We let $j = \sqrt{-1}$ then define $\zeta_i = e^{-2\pi j(i-1)/N}$ for $i = 1, \ldots, N$. Then the eigenvalues of A_0 are $\{\zeta_i\}$ and the eigenvalues of L_0 are

$$\lambda_i = 1 - \zeta_i = 1 - \cos\theta_i + j\sin\theta_i, \quad \theta_i = \frac{2\pi(i-1)}{N}, \quad i = 1, \ldots, N,$$

with corresponding eigenvectors:

$$v_i = \frac{1}{\sqrt{N}}\left(1, \zeta_i^{-1}, \zeta_i^{-2}, \ldots, \zeta_i^{-(N-1)}\right)^\top.$$

Hence we have

$$\Re(\lambda_i) = 1 - \cos\theta_i = 2\sin^2\Big(\frac{\theta_i}{2}\Big), \qquad \Im(\lambda_i) = \sin\theta_i.$$

As usual $\lambda_1 = 0$ corresponds to the consensus eigenvector 1_N, while the remaining eigenvalues appear in complex conjugate pairs.

5.1.5.1 Disagreement Dynamics

Thus for the disagreement dynamics, (13) yields:

$$\|G\|_2^2 = \sum_{i=2}^{N} \frac{1}{2\,\Re(\lambda_i)} = \sum_{m=1}^{N-1} \frac{1}{4\sin^2(\pi m/N)} = \frac{N^2-1}{12}.$$

Thus disagreement volatility grows like N^2, reflecting many slow modes. For worst-case amplification, from (15) only the smallest real part matters:

$$\|G\|_\infty = \frac{1}{\min_{i\geq 2}\Re(\lambda_i)} = \frac{1}{1-\cos(2\pi/N)} = \frac{1}{2\sin^2(\pi/N)} \sim \frac{N^2}{2\pi^2} \quad \text{as } N \to \infty.$$

Unlike symmetric graphs, the eigenvalues are complex, so the peak need not occur at frequency zero. As in (14) above, the worst-case frequency is attained at the mode with the smallest real part, which here is the conjugate pair $k = 2$ and $k = N$. For $k = 2$, $\Im(\lambda_2) = \sin(2\pi/N)$ and thus

$$\omega^* = -\Im(\lambda_2) = -\sin\Big(\frac{2\pi}{N}\Big),$$

while for $k = N$ we obtain the symmetric peak at $\omega^* = \sin(2\pi/N)$. Hence the peak occurs at $|\omega^*| = \sin(2\pi/N)$. For example, with $N = 4$ the eigenvalues are

$$\lambda_1 = 0, \qquad \lambda_2 = 1 - j, \qquad \lambda_3 = 2, \qquad \lambda_4 = 1 + j,$$

so $\|G\|_\infty = 1$ and the peak is at $\omega^* = 1$.

[28] See Section 3.1 and Theorem 3.1 of Gray (2006). Van Mieghem (2023) provides the related spectrum of an undirected cycle.

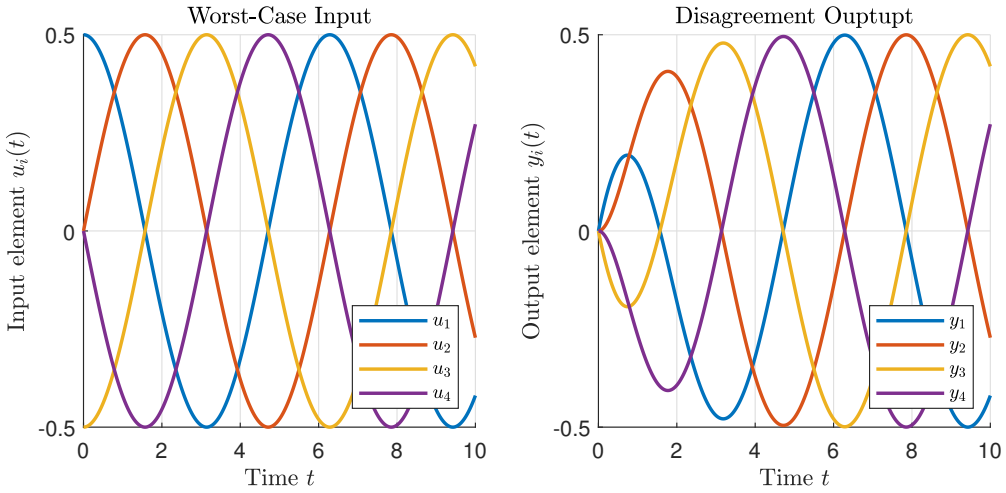


Figure 7: Worst-case input $u(t)$ and output $y(t)$ for disagreement dynamics on a 4-node directed cycle network.

Previously, the worst-case inputs were all essentially static, steady-state distortions, and hence were real vectors. However here, the worst-case input is dynamic. To visualize the worst-case disturbance, note that although the maximizing mode $\overline{v} = v_2$ is complex, the corresponding real-valued input can be written as

$$u(t) = \Re\big(e^{j\omega^\star t} v_2\big).$$

This produces sinusoidal waves that make each node oscillate at frequency ω^* but with a different time offset, producing a traveling wave around the cycle. For example, with $N = 4$ we have $v_2 = \frac{1}{2}(1, -j, -1, j)^\top$ and $\omega^\star = 1$, so

$$u(t) = \Re\big(e^{jt} v_2\big) = \frac{1}{2}\big(\cos t,\ \sin t,\ -\cos t,\ -\sin t\big)^\top.$$

This input is illustrated in the left panel of Figure 7. The worst-case for disagreement is to magnify differences at opposite parts of the cycle, and shift this disagreement over time. The components are quarter-period out of phase, and as t increases the pattern rotates around the cycle: at $t = 0$ the disturbance loads on nodes $(1, 3)$, at $t = \pi/2$ it loads on $(2, 4)$, and so on. The right panel of the figure plots the resulting disagreement vector $y(t)$, which starts at the steady-state of zero and then after a transient phase converges to the input $u(t)$, illustrating maximum unitary amplification in this example. This type of interior-frequency amplification and oscillatory dynamics did not arise in the undirected networks.

5.1.5.2 Production Networks: Cross-Section, Homogeneous Labor

We now turn to production networks on the directed cycle. With homogeneous labor shares the production Laplacian is again $L = (1 - \beta)I + \beta L_0$, and hence its eigenvalues are

$$\lambda_i = 1 - \beta\zeta_i = 1 - \beta\cos\theta_i + j\beta\sin\theta_i,$$

where ζ_i and θ_i are as defined above. As always, the $\mathcal{H}_\infty$ norm reduces to $1/(1-\beta)$. Unlike disagreement dynamics, directedness does not move the worst-case away from frequency zero, and the dominant mode is still the uniform mode. For the $\mathcal{H}_2$ norm, (11) yields:

$$\|G\|_2^2 = \frac{1}{2}\sum_{i=1}^{N}\frac{1}{\Re(\lambda_i)} = \frac{1}{2}\sum_{i=1}^{N}\frac{1}{1-\beta\cos\theta_i}.$$

This admits a closed form. Let $r = \beta/(1+\sqrt{1-\beta^2})$. Then

$$\sum_{i=1}^{N}\frac{1}{1-\beta\cos\theta_i} = \frac{N}{\sqrt{1-\beta^2}}\cdot\frac{1+r^N}{1-r^N}.$$

Since $r^N \to 0$ with N, we have:

$$\|G\|_2^2 = \frac{N}{2\sqrt{1-\beta^2}}\cdot\frac{1+r^N}{1-r^N} \sim \frac{N}{2\sqrt{1-\beta^2}} \text{ as } N\to\infty.$$

Cross-sectional volatility is driven by the density of slowly damped long-wavelength modes. It scales linearly with the number of sectors even though no sector is central.

5.1.5.3 Production Networks: Cross-Section, Heterogeneous Labor Shares

Unlike the earlier examples, the dynamics are no longer normal with heterogeneous labor shares. Under heterogeneous labor shares, the Laplacian takes the form $L = D - \beta A_0$ with $D = \text{diag}(d_1,\ldots,d_N)$ and $d_i = 1-\beta_i+\beta$. Such an L is no longer normal because while A_0 is normal it is not symmetric.[29] Thus with heterogeneity, we need to consider the more general formulas (7) and (4), which depend on more than just the spectrum of L.

Later I consider the implications of non-normality more broadly, but here I illustrate its effects in this example. I consider a simple case of heterogeneity where, as in the path example above, $N-1$ sectors have a common intermediate share β, while one sector has a different share β_1. In Figure 8, I set $N = 21$ and plot the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ norms, fixing $\beta = 0.6$ and varying β_1. Shown are the true values of the norms, along with the normal proxies calculated using the Laplacian eigenvalues. When β_1 is close to β, the normal proxies are close to the true values. They diverge particularly when β_1 increases toward one and thus becomes a relative bottleneck, and especially in the worst-case.

The right panel shows the worst-case input vector $\overline{v}$ when $\beta_1 = 0.95$ at location $i = 11$. Here this is the right singular vector associated with the largest singular value of $G(0) = L^{-1}$. For comparison, the figure also shows the eigenvector associated with the eigenvalue of L with the smallest real part, which would be the worst-case input if the dynamics were normal.

[29] Precisely, we have $LL^\top - L^\top L = \beta(A_0D - DA_0) + \beta(DA_0^\top - A_0^\top D)$. This vanishes only when A_0 commutes with D, which here requires $\beta_i = \beta$.

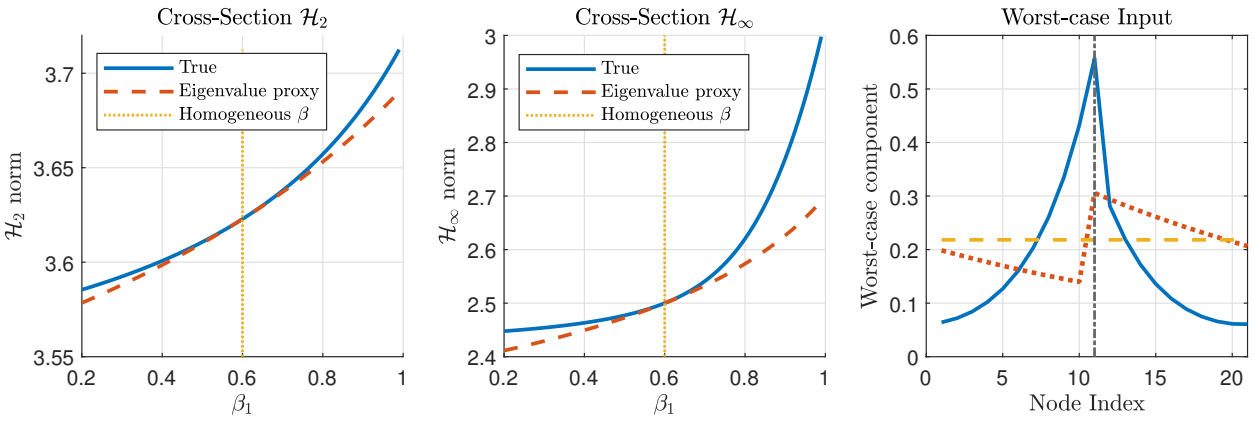


Figure 8: The left and center panels plot the $\mathcal{H}_2$ and $\mathcal{H}_\infty$ norms for a 21-node directed cycle production network with cross-sectional output, where 20 sectors have share parameter $\beta = 0.6$ and one sector has β_1. Shown are the true values along with the normal proxies calculated using the Laplacian eigenvalues. The right panel plots the worst-case input vector $\overline{v}$ when $\beta_1 = 0.95$.

Relative to the homogeneous case, the worst-case input now loads very significantly on the bottleneck sector, with weights that decay with distance from it. Moreover, since upstream and downstream effects are now distinct, the input vector is asymmetric. For example node 10 feeds directly into the bottleneck sector, while node 12 must pass through the whole cycle before it hits the bottleneck. Thus the worst-case input weights decay more rapidly downstream of the bottleneck than they do upstream. Note as well that the normal-proxy eigenvector gives a quite different input pattern.

To understand these effects, it is helpful to see how this cycle network operates on a steady-state input. Letting $x = L^{-1}u$, then $Lx = u$ implies the componentwise equations for each $i = 1, \ldots, N$ (modulo N),

$$d_i x_i - \beta x_{i-1} = u_i.$$

Rearranging yields the one-step recursion

$$x_i = \frac{\beta}{d_i}\, x_{i-1} + \frac{1}{d_i}\, u_i.$$

A shock u_i feeds directly into x_i with gain $1/d_i$, while upstream activity x_{i-1} is transmitted forward to x_i with gain β/d_i. Hence a sector with a small labor share, hence high β_i and small d_i, amplifies both its own disturbance and the upstream state. The directed cycle propagates an asymmetric downstream "wake" in the cross-sectional response. This helps explain why the worst-case input shifts toward the bottleneck sector and is asymmetric around it.

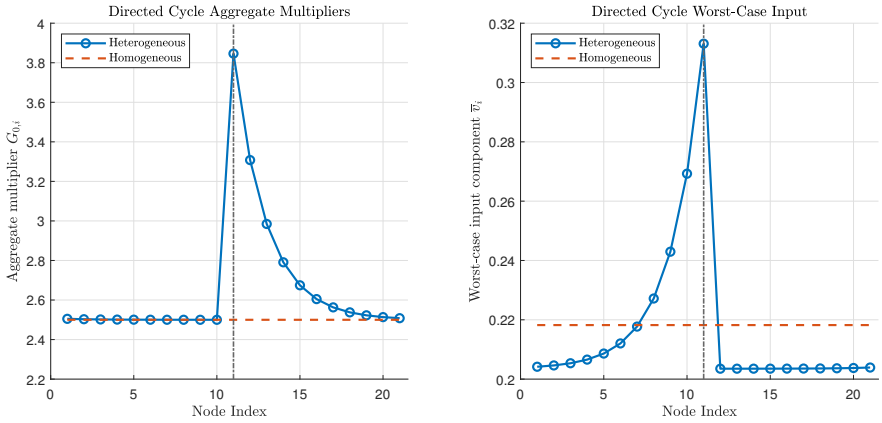


Figure 9: The left panel plots the aggregate multiplier vector G_0 and the right panel plots the worst-case input $\overline{v}$ for a 21-node directed cycle production network with aggregate output, where 20 sectors have $\beta = 0.6$ and one (node 11) has $\beta_1 = 0.95$. The homogeneous case is shown for comparison.

5.1.5.4 Production Networks: Aggregate Output

As above, with homogeneous labor shares the system norms collapse to expressions independent of the network structure, even though the graph is directed. As in the path graph, the aggregate multipliers under heterogeneity admit an explicit recursion. Writing out $LG_0 = 1_N$ gives a version of the recursion in the previous section (with $u_i \equiv 1$):

$$g_i = \frac{\beta}{d_i}\, g_{i-1} + \frac{1}{d_i}.$$

Let $a_i = \beta/d_i$, $c_i = 1/d_i$, and $\overline{a} = \prod_{i=1}^{N} a_i = \beta^N / \prod_{i=1}^{N} d_i$. Unwinding one full lap around the cycle yields the closed form

$$g_i = \frac{c_i + a_i c_{i-1} + a_i a_{i-1} c_{i-2} + \cdots + a_i a_{i-1} \cdots a_{i-N+2} c_{i-N+1}}{1 - \overline{a}},$$

with indices interpreted modulo N.

This recursion illustrates the directional propagation mechanism, as a low labor share sector (small d_i, hence large a_i) amplifies multipliers downstream around the cycle until the effect wraps around. The feedback term $1/(1 - \bar{a})$ captures the additional amplification from repeated laps. Further, the worst-case input is no longer proportional to the multiplier vector $G_0 = L^{-1}1_N$. Instead we have:

$$\overline{v} \propto \left(1_N^\top L^{-1}\right)^\top = L^{-\top} 1_N,$$

which only coincides with G_0 when $L = L^\top$. These effects are illustrated in Figure 9 for the same settings as in the cross-sectional case. The asymmetries from the directed

cycle are much more pronounced on aggregate output. The aggregate multipliers $G_{0,i}$ are highest on the bottleneck sector and those immediately downstream from it, which are in the downstream "wake" of the amplification. By contrast, the worst-case inputs are largest on the bottleneck sector and those immediately upstream, since upstream shocks feed most directly into the bottleneck and therefore move aggregate output the most.

5.2 Examples with Non-Normal Dynamics

The previous examples illustrated a number of features of network dynamics analytically, relying on the particular structure provided by normal dynamics. However once we allow for direction, asymmetry, and heterogeneity, normality is clearly an exceptional special case. Normality corresponds to network structures where influence patterns are symmetric or otherwise "balanced" in a spectral sense. The modes do not interfere, and eigenvalues fully characterize both shock propagation and worst-case amplification. As soon as the Laplacian governing the dynamics becomes non-normal, eigenvectors need not be orthogonal, modes interact, and small directional asymmetries in influence can create specific patterns of shocks, which are identified by the dominant singular vectors, that generate much larger amplification than the eigenvalues alone would predict.

In the normal case we were able to draw on known results for the spectra of classes of graphs, but there are many fewer analytic characterizations of singular value decompositions. In this section I focus on some small ($N = 2$ or 3) examples where I can provide analytic characterizations to highlight some of the features of non-normal dynamics. Beyond small cases, the results typically require numerical calculations, which however are straightforward.

5.2.1 Directed Two-Sector Production Network

We already saw an example of non-normality in the example (3) in Section 3.5. Here I analyze in more detail an expanded version of that example. Consider a simple 2-sector production network where neither sector uses its own good as an input, and they have different intermediate shares, $\beta_1 = a$ and $\beta_2 = b$, with each β_i strictly between zero and one. Then the adjacency matrix is

$$A \;=\; \begin{pmatrix} 0 & a \\ b & 0 \end{pmatrix},$$

and as long as $a \neq b$, A is non-normal and so is $L = I - A$.

Notice that this is a different, an more direct way, of handling heterogeneous input shares than I used above. Previously, I introduced share heterogeneity by scaling the diagonal terms of the adjacency, varying how much of its own input each sector used. Here I directly change

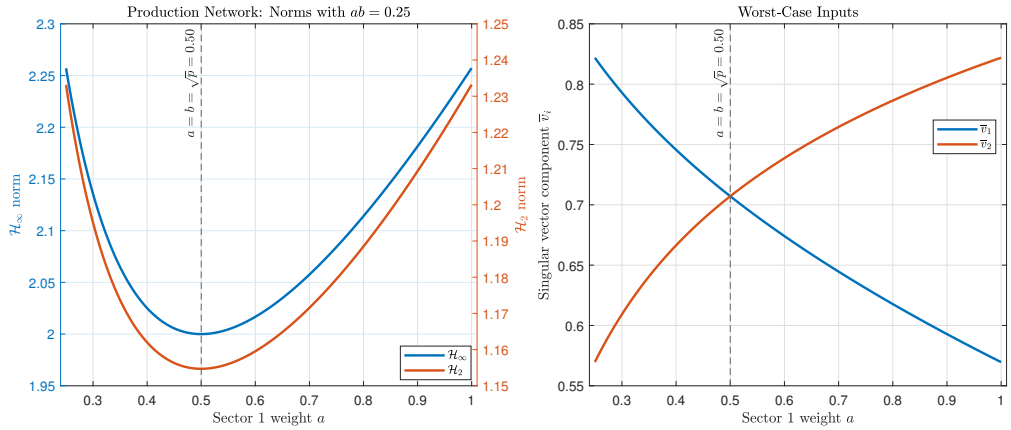


Figure 10: The left panel plots the $\mathcal{H}_\infty$ (left axis) and $\mathcal{H}_2$ (right axis) norms for a 2-sector directed production network. The sectors have input shares a and b with fixed product $ab = 0.25$, and the panel shows the norms for varying a. The right panel shows the corresponding worst-case input vector $\overline{v}$ as a varies.

the weight on other sectors. When $a < b$, sector 1 is relatively less-dependent on sector 2's good as an input, with sector 2 relatively more-dependent on sector 1.

The results of this section are illustrated in Figure 10, which shows the $\mathcal{H}_\infty$ and $\mathcal{H}_2$ norms for this 2-sector directed production network, focusing on cross-sectional output $C = I$. The figure fixes the product of the input shares at $ab = 0.25$ and shows the norms for varying a. The right panel plots the elements of the corresponding worst-case input vector $\overline{v}$ as a varies. The figure clearly shows that increasing sectoral heterogeneity increases both worst-case amplification and output volatility. Moreover, the worst-case input vector loads more heavily on the sector with the smaller input share coefficient. In Appendix B.2, I work out these results underlying the figure analytically, while here I interpret them.

In a two-sector production chain, a and b summarize how strongly each sector relies on the other. The eigenvalue of L reflects the average strength of the feedback loop $1 \leftrightarrow 2$ through the product ab. By contrast, the singular-vector emphasizes directionality. If one link is strong and the reverse link is weak, then disturbances can be transformed as they circulate. The operator $L^\top L$ aggregates these "out-and-back" effects, and its eigenvectors identify the disturbance patterns that are hardest for the network to dissipate.

When the sectors have different input intensities, the output volatility increases and the netowrk becomes more fragile. When $b > a$ sector 2 relies more heavily on sector 1 than vice versa. Sector 1 is therefore more upstream, as shocks to sector 1 propagate strongly to sector 2, while shocks to sector 2 feed back only weakly to sector 1. This directional imbalance shifts the worst-case amplification toward disturbances concentrated in sector 1, even when the overall loop strength ab, and hence the spectrum, is unchanged.

This two-node example illustrates several features of non-normality in directed production networks. First, there is a clear divergence between the eigenvalues and the singular values, with the latter governing the response of the network to exogenous shocks. Holding fixed the product ab fixes the eigenvalues, and hence common measures of the speed of convergence. But fixing ab while increasing the asymmetry between the two sectors increases both volatility and worst-case amplification of shocks. With this asymmetry, the worst-case inputs and most-amplified outputs differ when the network is directed. In this example, the difference is stark because an input which loads more heavily on one sector has the most impact on the output of the other. Systemic vulnerability reflects not only how strongly sectors are linked in aggregate, but how unevenly and directionally those linkages are organized, where the bottlenecks sit, and whether dependence is balanced across stages.

5.2.2 Three-Node Non-Normal Disagreement Dynamics

To further illustrate the effects of non-normality, I consider a variation on a directed cycle from Section 5.1.5. Here I introduce heterogeneity in the weights of the cycle, adding self-loops so each row sums to one. In particular, I consider a "lazy" directed 3-cycle, where two edge weights are a and one is b:

$$A_0 = \begin{pmatrix} 1-a & a & 0 \\ 0 & 1-a & a \\ b & 0 & 1-b \end{pmatrix},$$

where both a and b are in $(0, 1)$. When $a = b$, the dynamics are normal, but otherwise are non-normal. I focus here on disagreement dynamics where with $a \neq b$ the flow of information around the cycle is uneven, with one direction stronger than the reverse. Here I give an overview of the results, with explicit calculations in Appendix B.3.

Disagreement dynamics track deviations from the mean belief. To derive analytic results, I work on the two–dimensional mean–zero subspace with basis vectors

$$q_1 = \frac{1}{\sqrt{2}}(1, -1, 0)^\top, \qquad q_2 = \frac{1}{\sqrt{6}}(1, 1, -2)^\top.$$

Here q_1 is pairwise disagreement between nodes 1 and 2, while q_2 contrasts node 3 with the average of nodes 1 and 2. The two-dimensional disagreement $z(t)$ on the subspace satisfies $y(t) = [q_1, q_2]z(t)$. Since this transformation is isometric, both performance norms can be studied on it.

In the appendix I show that, with this change of variables, the $\mathcal{H}_2$ norm decomposes additively, with two terms corresponding to the two orthogonal disagreement channels (q_1, q_2):

$$\|G\|_2^2 = \frac{1}{3a} + \frac{1}{2b + a}.$$

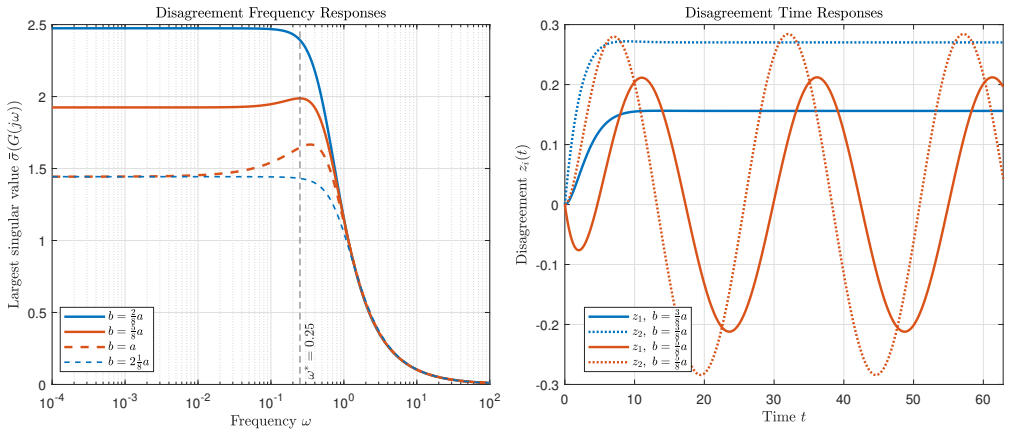


Figure 11: The left panel plots the maximal singular values for disagreement dynamics on a 3-node lazy directed cycle with fixed weight $a = 0.4$ and varying weight b. The right panel plots the time responses of the disagreement, measured in the orthonormal basis, for two of the cases from the right panel.

Thus heterogeneity in the directed cycle shows up as heterogeneity in mean-reversion speeds for distinct disagreement patterns. Shocks that look like a dispute between nodes 1 and 2 are dampened at an effective rate $3a$, while shocks that look like node 3 against the other two are dampened at rate $2b+a$. Even though all three nodes interact, average disagreement volatility is the sum of the contributions from these two orthogonal directions.

The worst-case amplification of disagreement splits into two different types of outcomes, as illustrated in Figure 11. For sufficiently small or sufficiently large b (relative to a), the minimizing frequency is zero. The left panel of the figure plots the largest singular value for fixed $a = 0.4$, where with b small (blue solid line) or large (blue dashed line) the peak is at zero. In this case, the $\mathcal{H}_\infty$ norm is determined by the weakest of the two restoring forces in the two orthogonal channels. The worst-case disturbance is essentially a persistent bias which pushes disagreement in a fixed cross-sectional direction. When b is small, the worst-case input loads onto q_2. Node 3 is weakly pulled toward the others, so the network is most vulnerable to disturbances that separate node 3 from the others. When b is large, the worst-case input loads on q_1, as the weakest channel becomes the direct disagreement between nodes 1 and 2. The right panel of Figure 11 shows the time response of disagreement under the worst-case input. The blue solid and dotted lines show the components of the reduced vector $z(t)$ in the small b case, which after a transient period settle into constant disagreement.

For an intermediate range of parameters, the minimizing frequency is interior. The red lines in the left panel of Figure 11 illustrate two such cases. As b increases, the overall amplification falls and an interior peak arises. The worst-case input is then a cyclical signal that exploits the directional structure of the network, creating amplification through timing

and phase. The worst-case disturbance rotates through the two disagreement patterns with a phase shift, as shown by the red lines in the right panel of Figure 11. Moreover, the relative weight b affects the amplitude and phase of the cyclical pattern. Thus asymmetry and heterogeneity of link strength change the patterns of disagreement and affect whether fragility comes from static bottlenecks or dynamic amplification.

6 Conclusion

This paper develops an approach to evaluate the performance and robustness of dynamics on networks using concepts from control theory. I study two common measures of performance: the $\mathcal{H}_2$ norm which summarizes average volatility generated by stochastic shocks, and the $\mathcal{H}_\infty$ norm which captures worst-case amplification and robustness. I develop frequency-domain methods that separate impacts driven by low-frequency disturbances due to persistent biases that accumulate, from vulnerabilities that arise at interior frequencies through directional propagation and feedback. I show how in normal or undirected networks components of the network spectrum govern the performance measures, while also illustrating how these measures change when networks are directed and flows are asymmetric. I apply these methods to analyze network dynamics in applications to production networks and local information transmission, but the methods extend beyond these settings.

I develop several examples that can be solved explicitly and analytically, illustrating how connectedness, redundancy, and speeds of adjustment affect the network dynamics. Across the examples, performance and robustness depend on the distribution of local "leakages" or restoring forces, where bottlenecks sit in the topology and on the choice of output measure. Cross-sectional versus aggregate output can weight network modes very differently, and disagreement dynamics may give rise to cyclical fluctuations. When networks are directed and non-normal, the dynamics separate upstream and downstream nodes, which generate large amplification even when eigenvalue-based measures suggest strong stability.

Several extensions are natural. Here I focus on some classes of networks that have explicit solutions, but the framework can provide performance measures and diagnostics for larger, real-world networks. As in the production networks literature, input-output tables provide a direct measurement of link intensities, allowing a quantitative assessment of which industries drive volatility versus worst-case fragility under alternative shock processes and aggregation schemes. In information propagation, large-scale directed graphs, such as social media networks, can be used to measure both propagation of information and the vulnerabilities of disagreement or polarization to targeted injections of content or misinformation.

Finally, as I mentioned in the introduction, this paper focuses on analysis rather than synthesis, in that I take as given the network structure. But the methods here could be

tools to study network control and design. One direction is to directly study decisions or policies that alter diffusion or routing so as to improve performance and robustness, trading off stabilization benefits against intervention costs. A complementary direction is network design, characterizing the benefits of adding or reweighting links in performance compared to the costs of link formation and maintenance. These problems connect the positive analysis of propagation to normative questions about optimal resilience and the architecture of economic and social networks.

References

Acemoglu, D. and P. D. Azar (2020). Endogenous production networks. *Econometrica 88*(1), 33–82.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica 80*(5), 1977–2016.

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011). Bayesian learning in social networks. *Review of Economic Studies 78*(4), 1201–1236.

Acemoglu, D. and A. Ozdaglar (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications 1*(1), 3–49.

Amir, G., I. Arieli, G. Ashkenazi-Golan, and R. Peretz (2025). Granular DeGroot dynamics: A model for robust naive learning in social networks. *Journal of Economic Theory 223*, 105952.

Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics 9*(4), 254–280.

Bala, V. and S. Goyal (1998). Learning from neighbours. *Review of Economic Studies 65*(3), 595–621.

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006). Who's who in networks. wanted: The key player. *Econometrica 74*(5), 1403–1417.

Banerjee, A. and O. Compte (2024). Consensus and disagreement: Information aggregation under (not so) naive learning. *Journal of Political Economy 132*(8), 2790–2829.

Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics 107*(3), 797–817.

Baqaee, D. (2018). Cascading failures in production networks. *Econometrica 86*(5), 1819–1838.

Baqaee, D. and E. Farhi (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem. *Econometrica 87*(3), 1155–1203.

Baqaee, D. and E. Farhi (2020). Productivity and welfare measurement with production networks. *Econometrica 88*(5), 2039–2066.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy 100*(5), 992–1026.

Bloch, F., M. O. Jackson, and P. Tebaldi (2023). Centrality measures in networks. *Social Choice and Welfare 61*(2), 413–453.

Boehm, C. E., A. Flaaen, and N. Pandalai-Nayar (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake. *Review of Economics and Statistics 101*(1), 60–75.

Bonacich, P. (1987). Power and centrality: A family of measures. *American Journal of Sociology 92*(5), 1170–1182.

Boyd, S., P. Diaconis, P. Parrilo, and L. Xiao (2005). Symmetry analysis of reversible markov chains. *Internet Mathematics 2*(1), 31–71.

Brock, W. A. and S. N. Durlauf (2005). Local robustness analysis: Theory and application. *Journal of Economic Dynamics and Control 29*(11), 2067–2092.

Brock, W. A., S. N. Durlauf, and G. Rondina (2013). Design limits and dynamic policy analysis. *Journal of Economic Dynamics and Control 37*(12), 2710–2728.

Brouwer, A. E. and W. H. Haemers (2012). *Spectra of Graphs*. Universitext. New York, NY: Springer.

Carvalho, V. M., M. Covarrubias, and G. Nuño (2025). Planning against disasters in dynamic production networks. Working paper.

Carvalho, V. M. and A. Tahbaz-Salehi (2019). Production networks: A primer. *Annual Review of Economics 11*, 635–663.

Cheng, C., X. Han, X. Tong, Y. Wu, and Y. Xing (2023). Degree-weighted DeGroot learning. arXiv preprint arXiv:2311.07010.

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

DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003). Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics 118*(3), 909–968.

Dupor, W. (1999). Aggregation and irrelevance in multi-sector models. *Journal of Monetary Economics 43*(2), 391–409.

Elliott, M., B. Golub, and M. V. Leduc (2022). Supply network formation and fragility. *American Economic Review 112*(8), 2701–2747.

Elliott, M. and M. O. Jackson (2026). Supply chain disruptions, the structure of production networks, and the impact of globalization. Working paper.

Ellison, G. and D. Fudenberg (1993). Rules of thumb for social learning. *Journal of Political Economy 101*(4), 612–643.

Ellison, G. and D. Fudenberg (1995). Word-of-mouth communication and social learning. *Quarterly Journal of Economics 110*(1), 93–125.

Foerster, A. T., P.-D. G. Sarte, and M. W. Watson (2011). Sectoral versus aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy 119*(1), 1–38.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*(3), 733–772.

Gale, D. and S. Kariv (2003). Bayesian learning in social networks. *Games and Economic Behavior 45*(2), 329–346.

Ghosh, A. (2021). Learning under bayesian vis-à-vis DeGroot updating. *Journal of Economic Behavior & Organization 188*, 967–986.

Golub, B. and M. O. Jackson (2010). Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics 2*(1), 112–149.

Golub, B. and M. O. Jackson (2012). How homophily affects learning and diffusion in networks. *Quarterly Journal of Economics 127*(3), 1287–1338.

Golub, B. and E. Sadler (2016). Learning in social networks. In Y. Bramoullé, A. Galeotti, and B. W. Rogers (Eds.), *The Oxford Handbook of the Economics of Networks*, pp. 505–542. Oxford: Oxford University Press.

Goyal, S. (2023). *Networks: An Economics Approach*. Cambridge, MA: MIT Press.

Gray, R. M. (2006). Toeplitz and circulant matrices: A review. *Foundations and Trends in Communications and Information Theory 2*(3), 155—239.

Hansen, L. P. and T. J. Sargent (2008). *Robustness*. Princeton, NJ: Princeton University Press.

Hansen, L. P. and T. J. Sargent (2022). Structured ambiguity and model misspecification. *Journal of Economic Theory 199*, 105165.

Horn, R. A. and C. R. Johnson (2013). *Matrix Analysis* (Second ed.). Cambridge: Cambridge University Press.

Horvath, M. (1998). Cyclicality and sectoral linkages: Aggregate fluctuations from sectoral shocks. *Review of Economic Dynamics 1*(4), 781–808.

Huo, Z., A. A. Levchenko, and N. Pandalai-Nayar (2025). International comovement in the global production network. *Review of Economic Studies 92*(1), 365—-403.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton, NJ: Princeton University Press.

Katz, L. (1953). A new status index derived from sociometric analysis. *Psychometrika 18*(1), 39–43.

Kopytov, A., B. Mishra, K. Nimark, and M. Taschereau-Dumouchel (2024). Endogenous production networks under supply chain uncertainty. *Econometrica 92*(5), 1621–1659.

Lancaster, P. and M. Tismenetsky (1985). *The Theory of Matrices* (2 ed.). Academic Press.

Liu, E. and A. Tsyvinski (2024). A dynamic model of input–output networks. *Review of Economic Studies 91*(6), 3608—-3644.

Long, J. B. J. and C. I. Plosser (1983). Real business cycles. *Journal of Political Economy 91*(1), 39–69.

Mahmoud, M. S., A. Oyedeji, and Y. Xia (2021). *Advanced Distributed Consensus for Multi-Agent Systems*. Singapore: Springer.

Newman, M. E. J. (2018). *Networks* (2 ed.). Oxford: Oxford University Press.

Onatski, A. and N. Williams (2003). Modeling model uncertainty. *Journal of the European Economic Association 1*(5), 1087–1122.

Pass, R. (2019). *A Course in Networks and Markets: Game-theoretic Models and Reasoning*. Cambridge, MA: MIT Press.

Rantzer, A. (2012). Distributed control of positive systems. *Proceedings of the IFAC World Congress 45*(26), 492–497.

Schaal, E. and M. Taschereau-Dumouchel (2025). Echoes and delays: Time-to-build in production networks. Working paper.

Skogestad, S. and I. Postlethwaite (2005). *Multivariable Feedback Control: Analysis and Design* (2 ed.). Chichester: John Wiley & Sons.

Taschereau-Dumouchel, M. (2025). Cascades and fluctuations in an economy with an endogenous production network. *The Review of Economic Studies*, forthcoming.

Van Mieghem, P. (2023). *Graph Spectra for Complex Networks* (2 ed.). Cambridge: Cambridge University Press.

Veerman, J. J. P. and R. Lyons (2020). A primer on Laplacian dynamics in directed graphs. *Nonlinear Phenomena in Complex Systems 23*(2), 199–216.

vom Lehn, C. and T. Winberry (2022). The investment network, sectoral comovement, and the changing U.S. business cycle. *The Quarterly Journal of Economics 137*(1), 387–433.

Yadykin, A. and A. Iskakov (2021). Modal analysis of the Sylvester equation via orthogonal decompositions. *Automatica 131*, 109748.

Zhou, K. and J. C. Doyle (1998). *Essentials of Robust Control*. Upper Saddle River, NJ: Prentice Hall.

Zhou, K., J. C. Doyle, and K. Glover (1996). *Robust and Optimal Control*. Upper Saddle River, NJ: Prentice Hall.

A Extensions and Proofs

A.1 Dampened Consensus Dynamics

Rather than focusing on disagreement, another approach to studying information propagation is to modify the Laplacian itself to introduce a dampening force, similar to what we saw in the production networks. This leads to the modified Laplacian

$$L^{\beta} = I - \beta A,$$

where $\beta < 1$ acts as a friction or discount parameter. In the DeGroot context, this introduces inertia or stickiness in agents' belief adjustment, reducing the influence of neighbors relative to one's own prior. This formulation generates a system with all eigenvalues having positive real parts, ensuring global stability even in levels. A similar structure appears in the production network, where labor adds a leakage term that stabilizes the system. I focus on disagreement dynamics in the consensus framework, but the analysis of dampened consensus dynamics follows the same approach as production networks.

A.2 Modeling Uncertainty

A general and flexible approach to modeling uncertainty is to represent misspecification as a feedback distortion of the effective disturbance entering the model. We begin with a baseline nominal model:

$$y(s) = G(s)u_0(s), \qquad G(s) = C(sI - M)^{-1}B, \tag{19}$$

where u_0 denotes the primitive shocks under the nominal model.[30] The true model is unknown and may differ from (19) in its parameters, structure, or measurement relationships. Rather than attempting to model these deviations explicitly, I represent their effect as a distortion of the disturbance entering the nominal dynamics. Specifically, I adopt the feedback representation given in (6). The unknown but stable and norm-bounded perturbation $\Delta(s)$ captures deviations from the nominal model, so the input $u(s) = u_0(s) + \Delta(s)y(s)$ is a combination of shocks to the nominal model $u_0(s)$ and the uncertainty $\Delta(s)$.

This formulation captures a wide class of possible deviations from the nominal model. In particular, several common sources of uncertainty can be embedded in, or bounded by, a representation of the form (6).

1. **Uncertainty in network structure or dynamics.** Suppose the true system matrix differs from the nominal one,

   $$M = M_0 + \Delta_M,$$

   where Δ_M is unknown but bounded. The true dynamics are

   $$\dot{x} = (M_0 + \Delta_M)x + Bu_0.$$

   This can be written exactly as

   $$\dot{x} = M_0x + Bu_0 + \Delta_M x.$$

   In the frequency domain, this corresponds to a feedback operator acting on the state x. But since x and y are related by the nominal dynamics, this state-feedback uncertainty

[30] The nominal model in control theory, for example Zhou, Doyle, and Glover (1996), is also called the approximating model in economics, as in Hansen and Sargent (2008).

can often be embedded, and always conservatively represented, in the form (6). Uncertainty in the network thus induces a feedback distortion of the effective disturbance entering the nominal model.

2. **Input-channel uncertainty.** Suppose that shocks are distorted before entering the system, so that the actual disturbance is

$$u = u_0 + \tilde{u},$$

where $\tilde{u}$ captures both mismeasurement of shocks and misspecification of how they propagate through the system. If the distortion depends on current or past outcomes, for example because agents infer shocks from observed activity, then the distortion can be modeled as $\tilde{u} = \Delta y$, leading directly to the feedback structure (6). Even when such dependence is not modeled explicitly, bounding the effect of input distortion by a norm constraint on Δ provides a conservative but tractable robustness measure.

3. **Output-side uncertainty.** If the reported output differs from the true output, for example due to aggregation bias or delayed measurement, then

$$\tilde{y} = y + \Delta_y y.$$

Since the effective disturbance entering decision rules or subsequent dynamics depends on observed outcomes, such output-side distortions again induce feedback effects. Under mild regularity conditions, these effects can be absorbed into a representation of the form (6). When an exact embedding is not possible, the feedback formulation still provides an upper bound on the induced distortion.

Taken together, these examples illustrate that a broad class of model uncertainties can be represented, exactly or conservatively, as norm-bounded feedback distortions of the effective disturbance. This representation allows robustness to be analyzed using the $\mathcal{H}_\infty$ norm of the nominal transfer function, without committing to a specific structural model of misspecification.

A.3 Proof of Theorem 1

This proof sketch of the Small Gain Theorem follows Zhou, Doyle, and Glover (1996).

Proof. Substituting the feedback relation (6) gives:

$$y(s) = G(s)\big(u_0(s) + \Delta(s)y(s)\big),$$

so that

$$\big(I - G(s)\Delta(s)\big)y(s) = G(s)u_0(s).$$

Define the loop gain $T(s) = G(s)\Delta(s)$. Since both G and Δ are stable, the feedback interconnection is internally stable if and only if $I - T(s)$ is invertible with no poles in the closed right half-plane. A sufficient condition is that the induced $\mathcal{H}_\infty$ norm satisfies $\|T\|_\infty < 1$. Using the submultiplicativity of the induced norm,

$$\|T\|_\infty = \|G\Delta\|_\infty \leq \|G\|_\infty \cdot \|\Delta\|_\infty < \gamma \cdot \frac{1}{\gamma} = 1.$$

Thus, if $\|G\|_\infty \leq \gamma$, the feedback interconnection is robustly stable for all Δ with $\|\Delta\|_\infty < 1/\gamma$.

For necessity, we argue by contrapositive. If $\|G\|_\infty > \gamma$, choose $\omega_\star$ such that $\bar{\sigma}\big(G(j\omega_\star)\big) > \gamma$ and let $u_\star, v_\star$ be the corresponding left and right singular vectors with unit norm. Define the rank-one matrix

$$\Delta_0 \;\equiv\; \frac{1}{\bar{\sigma}(G(j\omega^\star))}\, u^\star v^{\star H}, \qquad \text{so that} \qquad G(j\omega^\star)\Delta_0 v^\star = v^\star,$$

and hence $I - G(j\omega^\star)\Delta_0$ is singular. To extend this pointwise construction to a stable transfer function, let

$$\phi(s) \;\equiv\; \frac{s + j\omega^\star}{s - j\omega^\star},$$

which is stable and satisfies $\|\phi\|_\infty = 1$ and $\phi(j\omega^\star) = 1$. Define

$$\Delta(s) \;\equiv\; \Delta_0\, \phi(s).$$

Then $\Delta \in \mathcal{H}_\infty$ and

$$\|\Delta\|_\infty \;=\; \|\Delta_0\|\, \|\phi\|_\infty \;=\; \frac{1}{\bar{\sigma}(G(j\omega^\star))} \;<\; \frac{1}{\gamma}.$$

Moreover, $\Delta(j\omega^\star) = \Delta_0$, so $I - G(j\omega^\star)\Delta(j\omega^\star)$ is singular, and therefore the closed-loop transfer $(I - G\Delta)^{-1}$ is not stable. This establishes necessity, if $\|G\|_\infty > \gamma$, then by construction there is a stable $\Delta \in \mathcal{H}_\infty$ with $\|\Delta\|_\infty < 1/\gamma$ that destabilizes the feedback interconnection.

■

A.4 Proof of Theorem 2.

The proof here follows standard calculations.

Proof. When $u(t)$ is Gaussian white noise, we can define $P(t) = E(x(t)x(t)^H)$. Then we have the dynamics:

$$\dot{P} = MP + PM^H + BB^H,$$

and hence the Lyapunov equation gives the steady-state value of P. Since $y(t) = Cx(t)$, then $E(y(t)y(t)^H) = CPC^H$, and hence

$$\|G\|_2^2 = E(y(t)^H y(t)) = \mathrm{tr}(CPC^H).$$

■

A.5 Proof of Theorem 3

Proof. The expression for Y_{ij} is a special case of the block–modal representation of solutions to the Lyapunov equation in Yadykin and Iskakov (2021) for diagonalizable systems. The additional matrix S_e enters when computing the $\mathcal{H}_2$ norm via the trace identity. Explicitly, using the results above the $\mathcal{H}_2$ norm is

$$\|G\|_2^2 = \operatorname{tr}(CPC^H) = \operatorname{tr}\big(CVYV^HC^H\big) = \operatorname{tr}\big(S_eY\big),$$

where the last equality uses the cyclicality of the trace and the definition of S_e. Then expanding the terms in the trace we have:

$$\|G\|_2^2 = \operatorname{tr}\big(S_eY\big) = \sum_{i,j} S_{e,ji}Y_{ij} = \sum_{i,j} \frac{S_{e,ji}B_{e,ij}}{\lambda_i + \overline{\lambda_j}}.$$

■

A.6 Proof of Theorem 4

This proof follows Rantzer (2012), adapted to my setting.

Proof. Because $A \geq 0$ with row sums < 1, all eigenvalues of $L = I - A$ satisfy $\Re\lambda_i(L) > 0$ and $-L$ is Metzler. Hence e^{-Lt} is entry-wise nonnegative for all $t \geq 0$. With $B, C \geq 0$, the impulse response

$$h(t) \;=\; Ce^{-Lt}B$$

is also entry-wise nonnegative. The transfer function may be written as the Laplace transform

$$G(j\omega) = \int_0^\infty e^{-j\omega t}h(t)\,dt.$$

For each entry (k, ℓ) of $G(j\omega)$ we have

$$\big|G_{k\ell}(j\omega)\big| = \Big|\int_0^\infty e^{-j\omega t}\,h_{k\ell}(t)\,dt\Big| \leq \int_0^\infty |e^{-j\omega t}|\,h_{k\ell}(t)\,dt = \int_0^\infty h_{k\ell}(t)\,dt \;=\; G_{k\ell}(0).$$

Hence the matrices satisfy the entry-wise inequality

$$|G(j\omega)| \;\leq\; G(0).$$

For any matrix M, the spectral norm obeys $\|M\|_2 \leq \||M|\|_2$, and if $0 \leq X \leq Y$ entry-wise with $Y \geq 0$, then $\|X\|_2 \leq \|Y\|_2$. Applying these facts to $X = |G(j\omega)|$ and $Y = G(0)$ gives

$$\|G(j\omega)\|_2 \;\leq\; \big\||G(j\omega)|\big\|_2 \;\leq\; \|G(0)\|_2 \qquad \forall\,\omega \in \mathbb{R}.$$

Since equality holds at $\omega = 0$, we conclude

$$\sup_{\omega\in\mathbb{R}} \|G(j\omega)\|_2 = \|G(0)\|_2 = \Big\|\int_0^\infty h(t)\,dt\Big\|_2 = \|CL^{-1}B\|_2.$$

■

B Detailed Calculations for the Examples

B.1 Calculations for the Star Network

Here I calculate the full spectrum of the star network. The spectrum can be decomposed into two parts. First, we have the leaf-difference modes, as for any vector $x = (0, x_2, \ldots, x_N)^\top$ with $\sum_{j=2}^{N} x_j = 0$, we have $A_0 x = ax$. Since this subspace has dimension $N-2$, a is an eigenvalue with multiplicity $N-2$. Then, following Section 2.3 of Brouwer and Haemers (2012), we can make an equitable partition of the graph into $\{\{1\}, \{2, \ldots, N\}\}$. The quotient M_0 here uses an orthonormal basis $\{(1, 0, \ldots, 0)^\top, (0, 1, \ldots, 1)^\top / \sqrt{N-1}\}$. The remaining eigenvalues are those of the 2×2 quotient:

$$M_0 = \begin{pmatrix} 0 & \frac{1}{\sqrt{N-1}} \\ \frac{1}{\sqrt{N-1}} & a \end{pmatrix}.$$

Because A_0 is row-stochastic, 1 is an eigenvalue, hence the other root is $\mathrm{tr}(M_0) - 1 = a - 1 = -\frac{1}{N-1}$. Therefore we have the full spectrum of A_0, and hence that of $L_0 = I - A_0$ which is:

$$\lambda_1 = 0,\ \lambda_2 = \ldots = \lambda_{N-1} = \frac{1}{N-1},\ \lambda_N = 1 + \frac{1}{N-1}.$$

For production networks with heterogeneous input shares, the calculations shift slightly. As before, we can define the leaf-difference modes, but now they reflect the leakage from labor. For any vector $x = (0, x_2, \ldots, x_N)^\top$ with $\sum_{j=2}^{N} x_j = 0$, we now have that $a_\beta = 1 - \beta a = 1 - \beta + \beta/(N-1)$ is an eigenvalue of L with multiplicity $N-2$. Similarly, letting $d_c = 1 - \beta_c + \beta$, the quotient block for the Laplacian becomes:

$$M_\beta = \begin{pmatrix} 1 & -\frac{\beta}{\sqrt{N-1}} \\ -\frac{\beta}{\sqrt{N-1}} & a_\beta \end{pmatrix}.$$

Therefore the remaining two eigenvalues of L are the eigenvalues of M_β, which we can denote as λ_+ and λ_-, where in particular $\lambda_- < a_\beta$. To be explicit, the eigenvalues are $\lambda_\pm = \frac{1}{2}\left(d_c + a_\beta \pm \sqrt{(d_c - a_\beta)^2 + \frac{4\beta^2}{N-1}}\right)$. For $\beta > 0$ the off-diagonal coupling in M_β is nonzero, so the smaller eigenvalue is strictly below both diagonal entries. In particular, $a_\beta - \lambda_- = \frac{1}{2}\left(\sqrt{(d_c - a_\beta)^2 + \frac{4\beta^2}{N-1}} - (d_c - a_\beta)\right) > 0$.

For aggregate output, $b = (1, \sqrt{N-1})^\top$ are the coordinates of the unit vector $1_N^\top$ in the basis defining M_beta. Then solving $M_\beta G_0 = b$ on the orthonormal basis yields the values of the multipliers explicitly: $g_1 = (a_\beta + \beta) \,/\, \left(d_c a_\beta - \frac{\beta^2}{N-1}\right)$, $g_\ell = \left(d_c + \frac{\beta}{N-1}\right) / \left(d_c a_\beta - \frac{\beta^2}{N-1}\right)$.

B.2 Calculations from Section 5.2.1

In this example, the spectrum of L depends only on the product of the weights ab, while the singular values, and hence the volatility and amplification of shocks in the network, depends on a and b separately. The eigenvalues of L are $\lambda_i = 1 \pm \sqrt{ab}$, so the system is stable and the slowest eigenmode decay rate is $1 - \sqrt{ab}$. But worst-case amplification is controlled by the singular values. In particular, from Theorem 4 we have:

$$\|G\|_\infty = \bar{\sigma}(L^{-1}) = \frac{1}{\min_i \sigma_i(L)}.$$

A direct calculation gives:

$$L^\top L = \begin{pmatrix} 1+b^2 & -(a+b) \\ -(a+b) & 1+a^2 \end{pmatrix}.$$

The squared singular values of L are the eigenvalues of $L^\top L$, and standard calculations yield:

$$\min_i \sigma_i(L) = \sigma_1(L) = \sqrt{\frac{(2+a^2+b^2) - \sqrt{(a^2-b^2)^2 + \ 4(a+b)^2}}{2}}.$$

In the symmetric case $a = b$ this reduces to $\sigma_1(L) = 1 - a$, as before. But now heterogeneity in sectoral input shares decreases the singular value $\sigma_1(L)$ and therefore increases the $\mathcal{H}_\infty$ norm. In particular, say we fix ab but increase the divergence between a and b. Then we fix the determinant of $L^\top L$ but increase its trace, and therefore decrease its smaller eigenvalue. When the sectors have different input intensities, the network becomes more fragile.

The worst-case input vector $\overline{v}$ is the corresponding right singular vector of L^{-1}, and hence left singular vector of L, which solves $LL^\top \overline{v} = \sigma_1^2(L)\overline{v}$. This vector can be written:

$$\overline{v} \propto \begin{pmatrix} a+b \\ 1+a^2-\sigma_1^2(L) \end{pmatrix}.$$

The most-sensitive output pattern is the corresponding left singular vector of L^{-1}, which solves $L^\top L\ \overline{u} = \sigma_1^2(L)\overline{u}$. Given the simple structure here, these vectors are closely related. If we let $\Pi = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ denote the coordinate-swap matrix, a direct calculation shows:

$$LL^\top \ = \ \Pi\,(L^\top L)\,\Pi.$$

So left and right singular vectors are simple permutations of each other: $\overline{v} = \Pi\,\overline{u}$. When $a \neq b$, L is non-normal and $\overline{v} \neq \overline{u}$, so the most damaging shock pattern and the most amplified output pattern do not coincide. In particular, if the most-damaging input is concentrated in node 1, then the most-sensitive output is concentrated in node 2.

We can be even more explicit here. Letting $\bar{v} = (\bar{v}_1, \bar{v}_2)^\top$, define $s \equiv a + b > 0$ and $d \equiv b^2 - a^2$. Then from the expression above we have

$$\frac{\bar{v}_1}{\bar{v}_2} = \frac{a+b}{1+a^2-\sigma_1^2(L)} = \frac{2(a+b)}{\sqrt{d^2+4s^2}+a^2-b^2} = \frac{2s}{\sqrt{d^2+4s^2}-d},$$

where the second equality uses the expression for $\sigma_1^2(L)$. Then we have $\bar{v}_1/\bar{v}_2 > 1$ if:

$$2s > \sqrt{d^2+4s^2} - d \Leftrightarrow \sqrt{d^2+4s^2} < 2s+d \Leftrightarrow d^2+4s^2 < (2s+d)^2 = d^2+4s^2+4sd \Leftrightarrow 4sd > 0.$$

Thus $\bar{v}_1/\bar{v}_2 > 1$ if and only if $b > a$.

For the $\mathcal{H}_2$ norm, we first solve the Lyapunov equation from Theorem 2, which here reduces to

$$LP + PL^\top = I.$$

In particular we can express its diagonal elements as:

$$P_{11} = \frac{a^2 - ab + 2}{4(1-ab)}, \qquad P_{22} = \frac{b^2 - ab + 2}{4(1-ab)}.$$

Then from Theorem 2, the squared $\mathcal{H}_2$ norm is:

$$\|G\|_2^2 = \mathrm{tr}(P) = \frac{(a-b)^2 + 4}{4(1-ab)}.$$

Here we see directly that for fixed product ab, as the divergence between a and b increases, output volatility as measured by the $\mathcal{H}_2$ norm increases.

B.3 Calculations for Section 5.2.2

With the disagreement output $y = \check{C}x$, we track the differences of beliefs from the average. By construction, the disagreement vector sums to zero. In order to derive analytic expressions here, it will be useful to focus on the two dimensional disagreement subspace. In particular, let $\{q_1, q_2\}$ be an orthonormal basis for the mean-zero subspace:

$$q_1 = \frac{1}{\sqrt{2}}(1,-1,0)^\top, \qquad q_2 = \frac{1}{\sqrt{6}}(1,1,-2)^\top, \qquad Q \equiv [q_1\ q_2] \in \mathbb{R}^{3\times 2}.$$

Then $Q^\top Q = I_2$ and $QQ^\top = \check{C}$. Write the state as $x = Qz + \bar{x}\,1_3$, where $z \in \mathbb{R}^2$ and $\bar{x} \in \mathbb{R}$. Since $\check{C}x = QQ^\top x = Qz$, disagreement output is isometric to z:

$$y = \check{C}x = Qz, \qquad \|y\|_2 = \|z\|_2.$$

Therefore all of the norms we're interested in for y can be found from z. Projecting the dynamics onto the disagreement subspace yields

$$\dot{z}(t) = -L_d z(t) + Q^\top u(t), \qquad y(t) = Qz(t),$$

with reduced Laplacian

$$L_d \equiv Q^\top L Q = \begin{pmatrix} \frac{3a}{2} & -\frac{\sqrt{3}\,a}{2} \\ \frac{\sqrt{3}(2b+a)}{6} & b+\frac{a}{2} \end{pmatrix}.$$

The transfer function from u to disagreement output is therefore

$$G(s) = \check{C}(sI+L)^{-1} = Q\,(sI+L_d)^{-1}Q^\top.$$

In particular, to find the $\mathcal{H}_2$ norm, we first solve the Lyapunov equation for reduced disagreement system, which has a diagonal solution:

$$L_d^\top P + P L_d = I \;\Rightarrow P = \operatorname{diag}\Big(\frac{1}{3a},\; \frac{1}{2b+a}\Big).$$

Thus the $\mathcal{H}_2$ norm is:

$$\|G\|_2^2 = \operatorname{tr}(P) = \frac{1}{3a} + \frac{1}{2b+a}.$$

Although all three nodes are coupled, we can decompose the disagreement volatility into the two orthogonal disagreement channels (q_1, q_2). The first channel is the pairwise disagreement between nodes 1 and 2, which is damped at effective rate $3a$. The second channel contrasts node 3 with the other two, and this disagreement is damped at rate effective $2b+a$. Heterogeneity in the directed cycle shows up as heterogeneity in the mean-reversion speeds of distinct patterns of disagreement.

Worst-case amplification depends on the singular values, and the expression above gives:

$$\bar{\sigma}(G(j\omega)) = \bar{\sigma}\left(Q\,(sI+L_d)^{-1}Q^\top\right) = \bar{\sigma}\left((j\omega I + L_d)^{-1}\right) = \frac{1}{\min_i \sigma_i(j\omega I + L_d)},$$

where the second equality uses the fact that Q has orthonormal columns. For the $\mathcal{H}_\infty$ norm, we maximize this across frequencies, where now the worst-case frequency need not be tied directly to the imaginary part of an eigenvalue as in the normal case.

Since the complex matrix $m(\omega) = j\omega I + L_d$ is 2×2, we can analytically characterize its singular values for any frequency. In fact, the squared singular values are the eigenvalues of the Hermitian matrix $S(\omega) = m(\omega)^H m(\omega)$. Explicitly, the diagonal elements of S can then be written as $s_{11}(\omega) = \omega^2 + 3a^2$ and $s_{22}(\omega) = \omega^2 + \frac{(2b+a)^2}{3}$, and the smallest singular $\sigma_1(m(\omega))$ satisfies:

$$\sigma_1^2(m(\omega)) = \frac{s_{11}(\omega) + s_{22}(\omega) - \sqrt{\big(s_{11}(\omega) - s_{22}(\omega)\big)^2 + \frac{4\omega^2(b+2a)^2}{3}}}{2}.$$

The singular value thus depends on the ω only through its square, which we denote $w = \omega^2$. Moreover the leading term inside the square root is independent of frequency, and I denote its value as W:

$$W = s_{11}(\omega) - s_{22}(\omega) = 3a^2 - \frac{(2b+a)^2}{3}.$$

Then we can write the squared singular value as a function of w as:

$$f(w) = \sigma_1^2(m(\sqrt{w})) = w + \frac{1}{2}\left(3a^2 + \frac{(2b+a)^2}{3}\right) - \frac{1}{2}\sqrt{W^2 + \frac{4w(b+2a)^2}{3}}.$$

Differentiating gives

$$f'(w) = 1 - \frac{(b+2a)^2}{3} \cdot \frac{1}{\sqrt{W^2 + \frac{4w(b+2a)^2}{3}}}.$$

An interior minimizer $w^* > 0$ thus satisfies $f'(w^*) = 0$, giving

$$\sqrt{W^2 + \frac{4w^*(b+2a)^2}{3}} = \frac{(b+2a)^2}{3},$$

which implies

$$w^* = -a^2 + 3ab - \frac{5}{4}b^2 = \frac{5}{4}(2a-b)\left(b - \frac{2a}{5}\right).$$

Hence, for fixed a, an interior minimizing frequency exists if and only if

$$\frac{2a}{5} < b < 2a.$$

Outside this range, the minimum is attained at $\omega^* = 0$. I consider these two cases separately.

If $\omega^* = 0$, then

$$\sigma_1^2(0) = \frac{1}{2}\left(3a^2 + \frac{(2b+a)^2}{3} - \sqrt{\left(3a^2 - \frac{(2b+a)^2}{3}\right)^2}\right) = \min\left\{3a^2,\ \frac{(2b+a)^2}{3}\right\}.$$

In particular, for fixed a, we will choose different branches dependent on whether b is small or large. The smaller of the two effective static damping channels determines the worst-case gain. For an interior minimizer ω^*,the first-order condition yields:

$$\sigma_1^2(m(\omega^*)) = f(w^*) = \frac{3}{4}b(4a-b).$$

In summary, fixing a and varying b, we can pair the minimizing frequency and value of the $\mathcal{H}_\infty$ norm as:

$$\omega^* = \begin{cases} 0, & b \le \frac{2a}{5}, \\ \sqrt{-a^2 + 3ab - \frac{5}{4}b^2}, & \frac{2a}{5} < b < 2a, \\ 0, & b \ge 2a. \end{cases} \qquad \|G\|_\infty = \begin{cases} \frac{\sqrt{3}}{2b+a}, & b \le \frac{2a}{5}, \\ \frac{2}{\sqrt{3b(4a-b)}}, & \frac{2a}{5} < b < 2a, \\ \frac{1}{\sqrt{3}\,a}, & b \ge 2a. \end{cases}$$

We now explicitly calculate the worst-case input, which at the maximizing frequency ω^* is the right singular vector of $G(j\omega^*)$ associated with $\bar{\sigma}(G(j\omega^*))$. Working with the reduced

space as above, we can focus on mean zero inputs ξ satisfying $u = Q\xi$. The reduced map from ξ to disagreement is $(j\omega I + L_d)^{-1} = m(\omega)^{-1}$. So the worst case input $\overline{\xi}$ is the right singular vector of $m(\omega^*)^{-1}$, or equivalently, the left singular of $m(\omega^*)$. Thus we look for the eigenvector of the matrix

$$m(\omega^*)m(\omega^*)^H = \begin{pmatrix} s_{11}(\omega^*) & j\omega^* \dfrac{2a+b}{\sqrt{3}} \\ -j\omega^* \dfrac{2a+b}{\sqrt{3}} & s_{22}(\omega^*) \end{pmatrix}$$

corresponding to its smaller eigenvalue. The worst-case input in the original node space is then $\overline{v} = Q\,\overline{\xi}$.

We now work through the cases. When $\omega^* = 0$, the key matrix reduces to $m(0)m(0)^H = \mathrm{diag}(3a^2, (2b+a)^2/3)$. Hence when $b \leq \frac{2a}{5}$, then $\sigma_1^2(0) = \frac{(2b+a)^2}{3}$ and the minimizing eigenvector of $m(0)m(0)^H$ is $\overline{\xi} = (0,1)^\top$. Thus the worst-case input is

$$\overline{v} = Q(0,1)^\top = q_2 = \frac{1}{\sqrt{6}}(1,1,-2)^\top.$$

When $b \geq 2a$, then $\sigma_1^2(0) = 3a^2$, so the minimizing eigenvector is $\overline{\xi} = (1,0)^\top$. Hence

$$\overline{v} = Q(1,0)^\top = q_1 = \frac{1}{\sqrt{2}}(1,-1,0)^\top.$$

In these two cases, the worst-case input is constant.

When $\frac{2a}{5} < b < 2a$, the minimizing frequency is ω^* is the square root of the value of w^* given above. An un-normalized eigenvector of $m(\omega^*)m(\omega^*)^H$ associated with its smaller eigenvalue $\sigma_1^2(\omega^*)$ can then be written:

$$\overline{\xi} \propto \begin{pmatrix} -j\omega^* \frac{2a+b}{\sqrt{3}} \\ s_1 1(\omega^*) - \sigma_1^2(\omega^*) \end{pmatrix}.$$

To ease some notation, define κ^* from the first element of this vector and K as its norm:

$$\kappa^* = \frac{\omega^*(2a+b)}{\sqrt{3}}, \qquad K = \sqrt{(\kappa^*)^2 + \left(\frac{b^2-4a^2}{2}\right)^2}.$$

Then we can write the worst-case input in normalized fashion as:

$$\overline{\xi} = \frac{1}{K}\begin{pmatrix} -j\,\kappa^* \\ \frac{b^2-4a^2}{2} \end{pmatrix},$$

or in original coordinates:

$$\overline{v} = Q\overline{\xi} = \frac{1}{K}\Big(-j\,\kappa^* q_1 + \frac{b^2-4a^2}{2}\, q_2 \Big).$$

Converting this complex sinusoid to a real time-domain input gives:

$$\begin{aligned} u(t) &= \Re\big(e^{j\omega^\star t}\,\overline{v}\big) = \frac{\kappa^\star}{K}\,\sin(\omega^* t)\,q_1 + \frac{b^2-4a^2}{2K}\,\cos(\omega^* t)\,q_2 \\ &= \frac{\kappa^\star}{K}\,\sin(\omega^\star t)\,\frac{(1,-1,0)^\top}{\sqrt{2}} + \frac{b^2-4a^2}{2K}\,\cos(\omega^\star t)\,\frac{(1,1,-2)^\top}{\sqrt{6}}. \end{aligned}$$