Robust Intervention in Networks*

Daeyoung Jeong Tongseok Lim Euncheol Shin†

May 26, 2026

Abstract

We study robust intervention in networks where a decision maker (DM) allocates resources to guide outcomes toward a target, while an adversarial Nature selects the network's dependence structure within a moment-consistent uncertainty set to maximize the DM's loss. We characterize the unique robust intervention and show that Nature's worst-case response exhibits a row-wise rank-1 property: uncertainty concentrates along a single direction aligned with the DM's intervention vector, reducing high-dimensional distributional ambiguity to a sign-pattern-driven fixed point. Applying the framework to four network topologies—pandemic mitigation, financial stabilization, supply chain management, and faculty recruitment—robust intervention trades off mean influence against correlation risk, yielding qualitatively different prescriptions from standard models. We further show that cross-block links generate disproportionately large costs of uncertainty, as they propagate worst-case dependence through multiple network segments simultaneously.

Keywords: robust intervention; network uncertainty; adversarial Nature; correlation risk; robust optimization

1 Introduction

Central authorities routinely intervene in interconnected systems where outcomes depend on a network of interactions whose structure is only partially known. A health

*We have benefited from conversations with Wonki Cho, Andy Choi, Ben Golub, Sung-Ha Hwang, Donggyu Kim, Duk Gyoo Kim, Jeong-Yoo Kim, Youngwoo Koh, Matthew Kovach, Wei Li, Omer Tamuz, and participants at several faculty seminars and conferences, including Network Science in Economics 2025 at Stanford University and 24th Annual SAET Conference in Ischia. Author names are listed in alphabetical order, and all authors contributed equally to this paper.

†Daeyoung Jeong: School of Economics, Yonsei University, Seoul, Republic of Korea. Email: daeyoung.jeong@gmail.com; Tongseok Lim: Mitchell E. Daniels, Jr. School of Business, Purdue University, IN, USA. Email: lim336@purdue.edu; Euncheol Shin: KAIST College of Business, Seoul, Republic of Korea. Email: eshin.econ@kaist.ac.kr.

authority allocating vaccines, a central bank injecting liquidity, or a government subsidizing supply chain capacity must reason not only about direct effects, but also about how their interventions propagate through the network. Yet, the very dependencies that govern this propagation—contact patterns between regions, exposures between banks, or input-output dependencies between firms—are typically observed only with substantial uncertainty. While a growing literature has formalized the design of optimal interventions in networks (e.g., Galeotti et al., 2020, 2024; Jeong and Shin, 2024; Li and Tan, 2025; Parise and Ozdaglar, 2023; Sun et al., 2023), this body of work largely treats the network structure as known. How a decision maker should design interventions when the network's dependence structure is itself uncertain remains an open question.

This paper develops a framework for robust intervention in networks under structural uncertainty. We model the problem as a zero-sum game between a decision maker (DM) and an adversarial "Nature." The DM allocates resources to guide network outcomes toward a target, while Nature selects the dependence structure of the network within an uncertainty set to maximize the DM's loss. By analyzing the worst-case scenarios that emerge under this framework, we characterize the DM's optimal intervention strategy, which accounts for both the mean influence of the network and the risks introduced by uncertainty.

The importance of accounting for network uncertainty is best illustrated by concrete policy problems. During a pandemic, public health authorities allocate medical supplies across regions where true mobility patterns are observed with substantial noise. During a financial crisis, central banks inject liquidity into commercial banks, but the resulting interbank exposure patterns shift in unpredictable ways (Liang, 2018). Similar informational frictions plague the allocation of capacity subsidies in interdependent supply chains and the integration of new agents into expanding organizational networks. In all these cases, the planner may know the marginal moments of the network links from historical data, but cannot reliably estimate how shocks across different links are correlated. Under this asymmetry, robust intervention requires a fundamentally different approach: rather than simply maximizing influence by targeting central agents, the planner must guard against the worst-case correlation risk that Nature can inflict.

Our zero-sum game formalizes this trade-off through a quadratic objective function that captures the interaction between the intervention strategy and the network's uncertain structure. Nature chooses the dependence structure to maximize this quadratic term, amplifying uncertainty in the most detrimental way for the DM. Anticipating this worst-case response, the DM devises a strategy that remains effective regardless of how Nature manipulates the network's dependence structure (Theorem 1).

Our main result characterizes the DM's unique robust intervention and the corresponding worst-case network configuration chosen by Nature. We show that Nature's adversarial response exhibits a row-wise rank-1 structure. That is, for each receiver, uncertainty concentrates along a single direction whose sign pattern perfectly aligns with the DM's intervention, allowing Nature to maximally amplify the harm from correlation risk. This characterization reduces a high-dimensional ambiguity over the network's joint distribution to a highly tractable, sign-pattern-driven object, providing a sharp economic interpretation of worst-case dependence.

We apply the framework to four economically motivated network topologies, each illustrating a distinct policy problem. The d-regular network captures pandemic mitigation with symmetric contact patterns, where we show that the optimal intervention intensity is inverted-U shaped in network density—revealing a tension between the leverage benefits of connectivity and the correlation risk it amplifies (Proposition 1). The core-periphery network captures financial stabilization, where a robust central bank optimally allocates liquidity to peripheral institutions even when the core dominates in mean influence, contrasting with the bang-bang corner solutions of standard intervention models (Proposition 2). The directed bipartite network captures supply chain management, where upstream transmission uncertainty shifts the planner's optimal mix of subsidies toward direct downstream interventions (Proposition 3). Finally, the network expansion case captures faculty recruitment, illustrating how a planner hedges against partial information when integrating a new agent into an existing collaboration network (Proposition 4).

We further analyze the comparative statics of the model by evaluating the marginal value of robust intervention with respect to individual link variances (Proposition 5, Proposition 6). We demonstrate that this marginal value perfectly isolates the economic *cost of uncertainty*, which is driven entirely by the penalty of the naive DM's relative over-exposure to worst-case correlation. Furthermore, when a naive intervention fundamentally misinterprets the required direction of intervention, Nature's strategic flexibility widens the gap between the ex-ante and ex-post values of robustness, inflicting an even steeper penalty. We show how network architecture governs these costs: in core-periphery configurations, cross-block links generate a substantially larger marginal value of robust intervention than within-block links because they propagate Nature's uncertainty through both network segments simultaneously. This yields a direct policy implication: a regulator with limited measurement capacity must strictly prioritize estimating the variance of systemic, cross-block conduits, as they contribute disproportionately to the cost of uncertainty.[1]

Taken together, our findings highlight that effective interventions in uncertain networks must balance two competing forces: the benefit of targeting high-influence agents and the cost of correlation risk that Nature can inflict on the most influential directions. By characterizing the robust intervention and decomposing its costs, our framework provides an approach to decision-making in network environments under structural uncertainty.

1.1 Related Literature

The current paper is broadly related to three strands of literature: network interventions, strategic interactions in networks, and robust mechanism design.

Intervention in networks. In the recent growing literature on intervention in networks, there are two approaches: (i) intervening in the network structure (e.g., Della Lena, 2024; Sun et al., 2023) and (ii) intervening in incentives within a given network structure (e.g., Belhaj et al., 2023; Galeotti et al., 2020, 2024; Jeong and Shin, 2024; Li and Tan, 2025; Parise and Ozdaglar, 2023). Our paper relates to both strands of this literature. Specifically, we study how a DM can intervene in agents' strategic incentives while accounting for uncertainty in the realized network, which arises from a game-like interaction between the DM and an adversarial Nature. Under certain constraints, the network structure is determined by Nature, an adversarial player whose objective opposes that of the DM. Galeotti et al. (2024) also examine robust interventions to improve market efficiency in the context of oligopolistic market competition. Their notion of robustness focuses on achieving a certain property with high probability. In contrast, our approach emphasizes worst-case scenario optimization, aligning more closely with the robust mechanism design literature (e.g., Bergemann

[1]In Online Appendix OB, we extend the framework to incorporate *higher-order interactions*, which arise naturally in network games, supply chains, and contagion models. Using a second-order approximation, we show that the rank-1 property of the worst-case scenario persists even when network effects propagate through indirect pathways.

and Morris, 2005; Carroll, 2015).

Strategic interactions in networks. Intervention models are applicable to various frameworks, such as public goods games (e.g., Allouch, 2015; Bramoullé and Kranton, 2007; Galeotti and Goyal, 2009) and social learning models (e.g., DeGroot, 1974; DeMarzo et al., 2003; Golub and Jackson, 2010, 2012). One particular area of the literature related to our research focuses on network games with uncertainty. Previous studies (e.g., Chaudhuri et al., 2024; Galeotti et al., 2010; Shin, 2021) examine uncertainty on the agents' side, where agents make equilibrium decisions based on incomplete information about the underlying networks. In contrast, the current paper considers uncertainty on the DM's side; the DM, who intervenes in agents' behavior, has limited information about the underlying networks.

Robust mechanism design. Our study relates to the growing literature on robust mechanism design. Bergemann and Morris (2005) develop a robust implementation framework that achieves equilibrium under minimal assumptions about agents' knowledge. In a principal-agent model, Carroll (2015) extends this approach by relaxing distributional assumptions and designing mechanisms that perform optimally under worst-case scenarios of informational uncertainty. Similar worst-case scenario and max–min approaches have been extensively applied not only in principal-agent settings (e.g., Frankel, 2014; Garrett, 2014; Kambhampati, 2024) but also in auction settings (e.g., Brooks and Du, 2024, 2021; Che, 2022; He and Li, 2022; Koçyiğit et al., 2019) and industrial organization theory (e.g., Guo and Shmaya, 2025).

To the best of our knowledge, the current paper is the first to study the robustness of a DM's intervention when there is uncertainty about the relevant network structure. As in the previous robust mechanism design literature, we analyze the DM's intervention strategy in a network setting to ensure ex-ante optimal outcomes under worst-case scenarios. In our model, due to the quadratic objective function of the DM, the worst-case scenario is represented by the correlation structure, which is in line with, for example, Cremer and McLean (1988), He and Li (2022), Myerson (1981). In terms of modeling uncertainty, Che (2022) and Koçyiğit et al. (2019) consider distributional robustness; in contrast, the current model approaches uncertainty as a matrix completion problem.

2 Model

2.1 Environment and Robust Objective

Notation. Throughout the paper, vectors are columns. For a vector $\mathbf{a}$, a_i denotes its ith entry. For any $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, $\langle \mathbf{a}, \mathbf{b} \rangle$ and $\mathbf{a} \otimes \mathbf{b}$ denote the inner and outer products, respectively, and $\|\mathbf{a}\|$ is the Euclidean norm. For a matrix $\mathbf{A}$, $\mathbf{A}_i$ is its ith row (as a vector), $\mathbf{A}^j$ its jth column, $\mathbf{A}_{ij}$ the (i, j) entry, and $\mathbf{A}^\intercal$ the transpose. The operator norm is denoted by $\|\mathbf{A}\|$ and the Frobenius norm by $\|\mathbf{A}\|_\mathbf{F}$.[2] When clear from context, a vector is treated as a matrix; for instance, $\mathbf{a} \otimes \mathbf{b} = \mathbf{a}\mathbf{b}^\intercal$. Let PD^k and PSD^k denote the sets of all $k \times k$ real symmetric positive definite and positive semi-definite matrices, respectively. Without loss of generality, assume that agent $n + 1$ is the new agent joining the existing network of n agents.

Network, uncertainty, and intervention. Let $N = \{1, \ldots, n\}$ denote the set of agents. The (random) influence network is an $n \times n$ matrix $\mathbf{G} \in \mathbb{R}^{n \times n}$, where $\mathbf{G}_{ij}$ captures the effect of agent j's allocation on agent i's outcome. A positive (negative) $\mathbf{G}_{ij}$ indicates a beneficial (detrimental) effect of j on i. $\mathbf{G}$ may be symmetric, as in network games (e.g., Galeotti et al., 2020), or asymmetric, as in social learning (e.g., Jeong and Shin, 2024). The entries $\mathbf{G}_{ij}$ are correlated random variables, with mean $\mathbf{m}_{ij} = \mathrm{E}[\mathbf{G}_{ij}]$ and variance $\mathbf{v}_{ij}^2 = \mathrm{Var}(\mathbf{G}_{ij})$ for each $i, j \in N$.

There is a decision maker (DM) who allocates a vector of resources $\mathbf{x} \in \mathbb{R}^n$ across agents. The DM knows $\mathbf{m}_{ij}$ and $\mathbf{v}_{ij}^2$ for each i, j, but not the covariances across links within each row of $\mathbf{G}$.[3] For a given allocation $\mathbf{x}$, the realized outcome vector is $\mathbf{G}\mathbf{x}$.

We focus on this one-shot formulation ($\mathbf{G}\mathbf{x}$) to cleanly isolate the mechanics of worst-case correlation risk. However, this is not a fundamental limitation: in Online Appendix OB, we show that the core rank-1 property persists under higher-order interactions and long-run equilibria (e.g., the Leontief inverse).

Adversarial Nature and timing. The decision problem can be described as a three-stage game between the DM and an adversarial Nature.[4]

[2] Since $\|\mathbf{A}\|$ corresponds to the spectral norm, if $\mathbf{A}$ is symmetric and positive semi-definite, then $\|\mathbf{A}\| = \lambda_{\max} \leq \sum_{i=1}^n \lambda_i = \|\mathbf{A}\|_\mathbf{F}$, where $\lambda_{\max} \geq 0$ is the largest eigenvalue of $\mathbf{A}$, and $\lambda_i \geq 0$ is the ith largest eigenvalue of $\mathbf{A}$. Equality holds if and only if $\mathbf{A}$ is a rank-1 matrix.

[3] We microfound this ambiguity in terms of limited joint observations and latent-factor structures in Subsection 2.2.

[4] An equivalent interpretation is that the DM entertains a set of priors over $\mathbf{G}$ consistent with the

1. In the first stage, the DM chooses an allocation vector $\mathbf{x} \in \mathbb{R}^n$ to make the resulting outcome $\mathbf{Gx}$ close to a target vector $\mathbf{z} \in \mathbb{R}^n$, while accounting for the cost of deviating from a reference allocation vector $\mathbf{x}^0 \in \mathbb{R}^n$.

2. In the second stage, adversarial Nature selects a joint distribution (dependence structure) for $\mathbf{G}$, consistent with the DM's moment information, to maximize the DM's loss.

3. In the third stage, $\mathbf{G}$ is realized and the outcome vector $\mathbf{Gx}$ is determined.

Because the DM moves first, she anticipates Nature's worst-case response when selecting allocation vector $\mathbf{x}$.

Robust objective. The DM's loss has two components: a quadratic loss from deviations of $\mathbf{Gx}$ from the target $\mathbf{z}$, and a quadratic cost of deviating from the reference allocation $\mathbf{x}^0$. Let Π denote the set of distributions over $\mathbf{G}$ satisfying the marginal moment restrictions

$$\mathrm{E}_\pi[\mathbf{G}_{ij}] = \mathbf{m}_{ij} \quad \text{and} \quad \mathrm{Var}_\pi(\mathbf{G}_{ij}) = \mathbf{v}_{ij}^2 \qquad \text{for all } i, j \in N.$$

The DM's robust optimization problem is formulated as follows:

$$\min_{\mathbf{x}\in\mathbb{R}^n} \max_{\pi\in\Pi} \quad \frac{1}{2}\Big(\mathrm{E}_\pi\left[\|\mathbf{Gx}-\mathbf{z}\|^2\right] + \left\|\mathbf{C}^{\frac{1}{2}}(\mathbf{x}-\mathbf{x}^0)\right\|^2\Big), \tag{1}$$

where $\mathbf{C}$ is a symmetric positive semi-definite matrix capturing the cost of deviating from reference allocation $\mathbf{x}^0$.[5]

Let $\mathbf{m}_i = \mathrm{E}_\pi[\mathbf{G}_i^\top] \in \mathbb{R}^n$ collect the mean influences toward agent i, and let $\mathbf{U}_i = \mathbf{G}_i^\top - \mathbf{m}_i \in \mathbb{R}^n$ denote the corresponding deviation vector. Define $\mathbf{M}_i = \mathbf{m}_i \otimes \mathbf{m}_i$ and $\mathbf{B}_i = \mathrm{E}_\pi[\mathbf{U}_i\mathbf{U}_i^\top]$. In Appendix A, we show that the expected quadratic loss admits the following decomposition:

$$\mathrm{E}_\pi[\|\mathbf{Gx}-\mathbf{z}\|^2] = \langle \mathbf{x}, \mathbf{Mx}\rangle + \langle \mathbf{x}, \mathbf{Bx}\rangle - 2\langle \psi, \mathbf{x}\rangle + \|\mathbf{z}\|^2, \tag{2}$$

where $\mathbf{M} = \sum_{i=1}^n \mathbf{M}_i$ aggregates mean influences, $\mathbf{B} = \sum_{i=1}^n \mathbf{B}_i$ aggregates within-row covariance, and $\psi = \sum_{i=1}^n \psi_i$ with $\psi_i = z_i\mathbf{m}_i$.

Expression (2) decomposes the effect of changing $\mathbf{x}$ into three channels: a mean influence component $\langle \mathbf{x}, \mathbf{Mx}\rangle$, an uncertainty component $\langle \mathbf{x}, \mathbf{Bx}\rangle$, and a target-

available moment information and evaluates policies using a max–min criterion; see Ben-Tal et al. (2009) for related robust optimization formulations.

[5]Without loss of generality, we take $\mathbf{C}$ to be symmetric because the objective is quadratic. The factor $\frac{1}{2}$ is a conventional normalization for quadratic objectives (e.g., Galeotti et al., 2020, 2024).

alignment component $-2\langle\psi, \mathbf{x}\rangle$. In the main text, we combine this with the intervention cost $\|\mathbf{C}^{\frac{1}{2}}(\mathbf{x}-\mathbf{x}^0)\|^2$ to obtain a quadratic representation of the DM's robust optimization problem.

Standing properties. We maintain two properties throughout the analysis:

(A) **Property A (independent responsiveness).** The aggregated mean influence matrix $\mathbf{M} = \sum_{i=1}^{n} \mathbf{m}_i \otimes \mathbf{m}_i$ is full rank (equivalently, $\mathbf{M} \succ 0$).

(B) **Property B (non-negligence).** The DM's robust intervention $\mathbf{x}^*$ solving DM's problem (1) has no zero entries.

Property A requires that no agent's mean influence vector is a linear combination of others'. Mathematically, this property implies that the DM's objective is strictly convex in $\mathbf{x}$, which yields a unique robust intervention. Property B rules out corner solutions and will be used to obtain uniqueness of Nature's worst-case dependence structure. In Appendix B, we further provide primitive sufficient conditions and an equivalent characterization of Property B.

To fix ideas, this framework naturally accommodates a variety of economic environments. For instance, in pandemic mitigation, $\mathbf{x}$ represents medical resources and $\mathbf{B}$ captures ambiguous local outbreaks. In financial stabilization, $\mathbf{x}$ represents liquidity injections and $\mathbf{B}$ captures systemic funding shocks. In supply chain management, $\mathbf{x}$ represents capacity subsidies and $\mathbf{B}$ captures correlated production disruptions. Finally, in faculty recruitment, the framework models the uncertain collaborative impact of a new hire. We formally explore these applications and their corresponding network topologies in Subsection 3.2.

2.2 Uncertainty Set and Microfoundations

We now construct the uncertainty set $\mathcal{B}$ that disciplines Nature's adversarial move. Although Nature in principle chooses a joint distribution $\pi \in \Pi$, we show below that under the DM's quadratic loss the relevant ambiguity collapses to a within-row covariance object, motivating $\mathcal{B}$ as the natural reduced form.

Row-wise uncertainty. For each agent $i \in N$, let $\mathcal{B}_i$ denote the set of matrices $\mathbf{B}_i$ that are symmetric and positive semi-definite,[6] and satisfy $(\mathbf{B}_i)_{jj} = \mathbf{v}_{ij}^2$ for all $j \in N$.

[6]This restriction is consistent with $\mathbf{B}_i = \mathrm{E}[\mathbf{U}_i\mathbf{U}_i^\mathsf{T}]$, where $\mathbf{U}_i$ is the vector of deviations in $\mathbf{G}_i^\mathsf{T}$.

The set $\mathcal{B}_i$, referred to as the *individual uncertainty set*, is convex and compact.[7] The aggregate uncertainty set is the Minkowski sum

$$\mathcal{B} \;=\; \Big\{\mathbf{B} \in \mathbb{R}^{n\times n} \;\Big|\; \mathbf{B} = \sum_{i=1}^{n} \mathbf{B}_i \text{ for some } \mathbf{B}_i \in \mathcal{B}_i\Big\},$$

which collects all feasible aggregated covariance objects Nature can choose. We refer to $\mathcal{B}$ as the *uncertainty set* in the sense of robust optimization (Ben-Tal et al., 2009).[8]

Moment-based ambiguity and row-wise dependence. We interpret $\mathcal{B}$ as arising from limited statistical information about the influence network. The random influence matrix is $\mathbf{G} \in \mathbb{R}^{n\times n}$, with entries $\mathbf{G}_{ij}$ capturing the effect of agent j on agent i. The DM observes a finite history of network realizations or reduced-form estimates and can reliably estimate, for each link (i,j), the mean $\mathbf{m}_{ij}$ and variance $\mathbf{v}_{ij}^2$ of $\mathbf{G}_{ij}$, but has only coarse information about dependence across links within each row of $\mathbf{G}$ because joint observations of $(\mathbf{G}_{ij}, \mathbf{G}_{ik})$ are limited.[9] Accordingly, the DM imposes the moment conditions $\mathrm{E}[\mathbf{G}_{ij}] = \mathbf{m}_{ij}$ and $\mathrm{Var}(\mathbf{G}_{ij}) = \mathbf{v}_{ij}^2$ for all $i,j \in N$, and treats as feasible every dependence pattern consistent with these constraints and covariance feasibility.

Writing $\mathbf{G} = \mathrm{E}[\mathbf{G}] + \mathbf{U}$ with $\mathbf{U}$ mean-zero, and letting $\mathbf{U}_i = \mathbf{G}_i^{\mathsf{T}} - \mathbf{m}_i \in \mathbb{R}^n$ denote the column vector of row-i deviations, the covariance matrix $\mathbf{B}_i = \mathrm{E}[\mathbf{U}_i\mathbf{U}_i^{\mathsf{T}}] \in \mathrm{PSD}^n$ has diagonal $(\mathbf{B}_i)_{jj} = \mathbf{v}_{ij}^2$ pinned down by the maintained moments, while the off-diagonal entries $(\mathbf{B}_i)_{jk} = \mathrm{Cov}(\mathbf{G}_{ij}, \mathbf{G}_{ik})$ for $j \neq k$ are left unspecified.

Correlation-matrix representation. Let $\mathbf{D}_i = \mathrm{diag}(\mathbf{v}_{i1}, \ldots, \mathbf{v}_{in})$. Then, any $\mathbf{B}_i \in \mathcal{B}_i$ can be written as $\mathbf{B}_i = \mathbf{D}_i\mathbf{R}_i\mathbf{D}_i$, where $\mathbf{R}_i$ is a correlation matrix (an element of the elliptope). With this decomposition of $\mathbf{B}_i$, it follows that uncertainty is entirely about within-row correlations.

Why $\mathcal{B}$ is the relevant reduced form under quadratic loss. As shown in Ap-

[7] Convexity is immediate. Compactness follows because $\mathcal{B}_i$ is closed and bounded: for any given $\mathbf{B}_i \succeq 0$, the Cauchy–Schwarz inequality implies $|(\mathbf{B}_i)_{jk}| \leq \sqrt{(\mathbf{B}_i)_{jj}(\mathbf{B}_i)_{kk}} = \mathbf{v}_{ij}\mathbf{v}_{ik}$ for all $j,k \in N$. In particular, $\mathrm{tr}(\mathbf{B}_i) = \sum_{j=1}^{n}(\mathbf{B}_i)_{jj} = \sum_{j=1}^{n}\mathbf{v}_{ij}^2$ for all $\mathbf{B}_i \in \mathcal{B}_i$.

[8] Our uncertainty concerns dependence in the influence network rather than payoff or technology uncertainty, as in robust mechanism design and related models; see, for example, Bergemann and Morris (2005); Carroll (2015); He and Li (2022).

[9] This interpretation parallels moment-based ambiguity sets in distributionally robust optimization, where the decision maker trusts marginal moments but remains agnostic about the joint distribution beyond these moments; see, for example, Delage and Ye (2010); Zymler et al. (2013).

pendix A, the expected quadratic loss depends on the distribution of $\mathbf{G}$ only through first and second moments. In particular, dependence ambiguity affects the objective only through the row-wise covariance objects $\{\mathbf{B}_i\}_{i\in N}$, hence through their aggregate $\mathbf{B} = \sum_{i=1}^{n} \mathbf{B}_i \in \mathcal{B}$.[10] This is why Nature's move can be represented as a choice of $\mathbf{B} \in \mathcal{B}$ rather than a choice over the full set of joint distributions for $\mathbf{G}$.[11]

3 Analysis

This section characterizes the DM's robust intervention and Nature's rank-1 worst-case response in a fixed network (Subsection 3.1), and applies these findings to four distinct network topologies that capture the structural features of our motivating economic applications (Subsection 3.2).

3.1 Main Theorem: Robust Intervention

This subsection presents the main characterization theorem for the DM's robust intervention and Nature's worst-case response in a fixed network. We work with the quadratic representation obtained in Appendix A. Let $\psi^0 = \mathbf{C}\mathbf{x}^0$ and recall that $\mathbf{M} = \sum_{i=1}^{n} \mathbf{m}_i \otimes \mathbf{m}_i$. For $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{B} \in \mathcal{B}$, we define

$$f(\mathbf{x}, \mathbf{B}) = \frac{1}{2}\Big(\langle \mathbf{x}, (\mathbf{M} + \mathbf{C})\mathbf{x}\rangle + \langle \mathbf{x}, \mathbf{B}\mathbf{x}\rangle - 2\langle \psi^0 + \psi, \mathbf{x}\rangle\Big), \tag{3}$$

where constants that do not depend on $(\mathbf{x}, \mathbf{B})$ are dropped. Then, the DM's robust problem (1) is equivalent to $\min_{\mathbf{x}\in\mathbb{R}^n} \max_{\mathbf{B}\in\mathcal{B}} f(\mathbf{x}, \mathbf{B})$.

Theorem 1 *Assume Properties A and B. Then, the following hold:*

(a) (Duality, saddle point, and robust intervention.) The minimax equality holds:

$$\min_{\mathbf{x}\in\mathbb{R}^n} \max_{\mathbf{B}\in\mathcal{B}} f(\mathbf{x}, \mathbf{B}) \;=\; \max_{\mathbf{B}\in\mathcal{B}} \min_{\mathbf{x}\in\mathbb{R}^n} f(\mathbf{x}, \mathbf{B}). \tag{4}$$

Moreover, there exists a saddle point $(\mathbf{x}^*, \mathbf{B}^*) \in \mathbb{R}^n \times \mathcal{B}$ such that $\mathbf{x}^*$ solves the DM's primal problem and $\mathbf{B}^*$ solves Nature's dual problem, and they satisfy the

[10] Online Appendix OC provides additional microfoundations for $\mathcal{B}$, including Bayesian and partial-identification interpretations and concrete data environments that rationalize ambiguity about within-row covariances.

[11] See also Gupta and Kallus (2022) for related Bayesian interpretations of moment-based ambiguity sets under quadratic loss.

first-order condition

$$(\mathbf{M} + \mathbf{B}^* + \mathbf{C})\mathbf{x}^* \;=\; \psi^0 + \psi. \tag{5}$$

The robust intervention $\mathbf{x}^*$ *is unique.*

(b) *(Rank-1 structure of the worst-case uncertainty.) Fix any* $\mathbf{x} \in \mathbb{R}^n$ *with no zero entries, and define, for each* $i \in N$*, the vector*

$$\mathbf{q}_i(\mathbf{x}) = \big(s(x_1)\mathbf{v}_{i1},\, \ldots,\, s(x_n)\mathbf{v}_{in}\big)^{\mathsf{T}} \in \mathbb{R}^n. \tag{6}$$

Then Nature's (row-wise) best response to $\mathbf{x}$ *is unique and given by* $\mathbf{B}_{i,\mathrm{BR}}(\mathbf{x}) = \mathbf{q}_i(\mathbf{x}) \otimes \mathbf{q}_i(\mathbf{x})$ *for each* $i \in N$*, and the aggregated best response is*

$$\mathbf{B}_{\mathrm{BR}}(\mathbf{x}) = \sum_{i=1}^{n} \mathbf{q}_i(\mathbf{x}) \otimes \mathbf{q}_i(\mathbf{x}) \in \mathcal{B}. \tag{7}$$

At the saddle point, $\mathbf{x}^*$ *has no zero entries by Property B, so the formula above applies and* $\mathbf{B}^* = \mathbf{B}_{\mathrm{BR}}(\mathbf{x}^*)$*. Consequently,* $\mathbf{B}^* = \sum_{i=1}^{n} \mathbf{B}_i^*$ *with* $\mathbf{B}_i^* = \mathbf{q}_i(\mathbf{x}^*) \otimes \mathbf{q}_i(\mathbf{x}^*)$ *for all* $i \in N$*, and the robust intervention satisfies the fixed-point equation:*

$$\mathbf{x}^* \;=\; \Big(\mathbf{M} + \mathbf{C} + \sum_{i=1}^{n} \mathbf{q}_i(\mathbf{x}^*) \otimes \mathbf{q}_i(\mathbf{x}^*)\Big)^{-1}(\psi^0 + \psi). \tag{8}$$

(c) *(Uniqueness of the worst-case scenario.) The worst-case uncertainty* $\mathbf{B}^*$ *is unique because* $\mathbf{x}^*$ *has no zero entries.*[12]

Computation via orthants. Theorem 1-(b) implies that Nature's response depends on $\mathbf{x}$ only through the sign pattern of its entries. For any sign vector $\mathbf{s} \in \{+1, -1\}^n$, define $\mathbf{q}_i(\mathbf{s}) = (s_1\mathbf{v}_{i1}, \ldots, s_n\mathbf{v}_{in})^{\mathsf{T}}$ and $\mathbf{B}(\mathbf{s}) = \sum_{i=1}^{n} \mathbf{q}_i(\mathbf{s}) \otimes \mathbf{q}_i(\mathbf{s})$. Then $\mathbf{x}(\mathbf{s}) = (\mathbf{M} + \mathbf{C} + \mathbf{B}(\mathbf{s}))^{-1}(\psi^0 + \psi)$ is well defined under Property A. A sign vector $\mathbf{s}$ is consistent if $\mathbf{s}$ matches the sign pattern of $\mathbf{x}(\mathbf{s})$, in which case $\mathbf{x}^* = \mathbf{x}(\mathbf{s})$ and $\mathbf{B}^* = \mathbf{B}(\mathbf{s})$.

Proof sketch of Theorem 1. By Appendix A, the robust objective reduces to $\min_{\mathbf{x}} \max_{\mathbf{B} \in \mathcal{B}} f(\mathbf{x}, \mathbf{B})$ with f quadratic in $\mathbf{x}$ and linear in $\mathbf{B}$. Under Property A, $f(\cdot, \mathbf{B})$ is strictly convex and coercive, while $f(\mathbf{x}, \cdot)$ is linear on the convex compact set $\mathcal{B}$. Sion's minimax theorem yields the minimax equality and a saddle point $(\mathbf{x}^*, \mathbf{B}^*)$, with $\mathbf{x}^*$ unique by strict convexity. Nature's problem separates across rows: writing $\mathbf{B}_i = \mathbf{D}_i\mathbf{R}_i\mathbf{D}_i$ with $\mathbf{D}_i = \mathrm{diag}(\mathbf{v}_{i1}, \ldots, \mathbf{v}_{in})$ and $\mathbf{R}_i$ a correlation matrix, $\langle \mathbf{x}, \mathbf{B}_i\mathbf{x} \rangle$ is maximized by aligning $\mathbf{R}_i$ with the sign pattern of $\mathbf{x}$, i.e., $(\mathbf{R}_i)_{jk} = s(x_j)s(x_k)$, yielding

[12] If Property B does not hold, then $\mathbf{B}_{\mathrm{BR}}(\mathbf{x}^*)$ need not be unique.

the rank-1 form $\mathbf{B}_i^* = \mathbf{q}_i(\mathbf{x}) \otimes \mathbf{q}_i(\mathbf{x})$. Property B ensures $\mathbf{x}^*$ is interior, pinning down the sign pattern and making Nature's best response unique.

3.2 Applications

To illustrate the flexibility and economic significance of our robust intervention framework, we now apply our analytical results to four distinct network topologies. Table 1 summarizes the mapping of our model primitives to these economic contexts.

Context	**Intervention (x)**	**Mean influence (M)**	**Uncertainty (B)**
Pandemic	Resource allocation	Contact and mobility patterns	Ambiguous local outbreaks
Finance	Liquidity injection	Interbank lending and exposure	Systemic funding shocks
Supply chain	Capacity subsidies	Input-output dependencies	Correlated production disruptions
Faculty recruitment	Research funding/support	Collaboration and citations	Unknown cross-link effects

Table 1: Mapping of Model Primitives to Economic Contexts

For each network topology, we first frame the mathematical structure within its motivating economic context, perform the robust optimization analysis in a context-free manner, and then interpret the optimal intervention strategy through the lens of the application. Each application below exhibits a permutation symmetry of $(\mathbf{M}, \mathbf{B}, \mathbf{C}, \mathbf{z}, \mathbf{x}^0)$, so the unique robust intervention inherits this symmetry and the n-dimensional problem reduces to a low-dimensional representative-agent problem (e.g., Galeotti et al., 2020; Jeong and Shin, 2024). Property A may fail for the ambient $\mathbf{M}$ (e.g., periphery agents share rows; upstream agents have zero rows), but holds for the reduced mean-influence matrix, which is the level at which Theorem 1 is applied throughout this subsection.

3.2.1 The d-Regular Network: Pandemic Mitigation

To analyze the robust allocation of medical resources (e.g., vaccines or testing kits) across a decentralized geographic area, we model the regions as a population of n agents arranged in a d-regular network G_d with $2 \le d \le n-1$ (e.g., Watts and Strogatz, 1998). Each agent i (a region) connects to exactly d neighbors in a symmetric manner, representing contact and mobility patterns. Figure 1 illustrates how the network topology varies with the density parameter d.

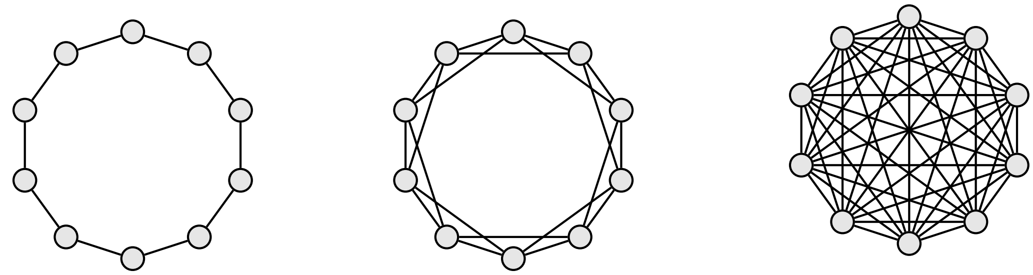

(a) sparse ring ($d = 2$) (b) dense lattice ($d = 4$) (c) complete ($d = 9$)

Figure 1: Illustration of d-regular networks of $n = 10$ agents.

We assume that each link transmits influence with a strictly positive mean influence $m > 0$ and a variance $v^2 \geq 0$. Due to the symmetry of the network, we restrict our attention to a symmetric intervention without loss of generality; that is, $x_i = x$ for all agents i in the network. Similarly, we assume that the intervention cost parameter is $c > 0$ and the target is $z_i = 1$ for all i. Consequently, by taking into account the worst-case scenario, the DM (e.g., a public health planner) minimizes the following objective function:

$$\frac{n}{2} \underbrace{(d^2\Omega + c)x^2}_{\text{quadratic penalty}} - \underbrace{ndmx}_{\text{leverage benefit}}, \tag{9}$$

where $\Omega = m^2 + v^2 > 0$ represents the aggregated link weight parameter. By the first-order condition, the optimal symmetric intervention intensity is given in Proposition 1:

Proposition 1 *The unique robust symmetric intervention is* $x^*(d) = \frac{dm}{d^2(m^2+v^2)+c}$. $x^*(d)$ *is inverted-U shaped and peaks at* $d^* = \sqrt{c/(m^2 + v^2)}$, *which is decreasing in* v.

For $d < d^*$, intervention increases with density due to the dominance of leverage; for $d > d^*$, it decreases as the correlation penalty dominates. Higher uncertainty lowers this threshold. This implies that optimal networks balance leverage and risk: sparse networks underutilize influence, whereas dense networks over-correlate shocks.

Proposition 1 offers a distinct view of the "robust-yet-fragile" property of dense networks (Acemoglu et al., 2015). In the context of pandemic policy, increasing mobility links (d) improves the average sharing of medical resources. However, it also creates *ex-ante informational fragility*: as d increases, the worst-case correlation risk of regional outbreaks grows in proportion to d^2, dominating the linear leverage benefit in d. Consequently, a robust public health planner facing highly connected, highly uncertain regions might actually reduce uniform intervention intensity, leading not to mechanical failure, but to an optimal shift toward more cautious, targeted strategies.

3.2.2 The Core-Periphery Network: Financial Stabilization

We next analyze the financial stabilization problem, mapping the interbank lending market to a core-periphery network (CP network). Denoted by $CP_{n,k}$, the network consists of n agents, where agents $1, \ldots, k$ form the "money-center" core and agents $k+1, \ldots, n$ form the regional periphery (e.g., Borgatti and Everett, 1999). As illustrated in Figure 2, the CP network is characterized by a dense core block (clique) and a periphery that connects only to the core. This core-periphery architecture is widely documented in interbank lending networks (e.g., Craig and von Peter, 2014; Fricke and Lux, 2015).

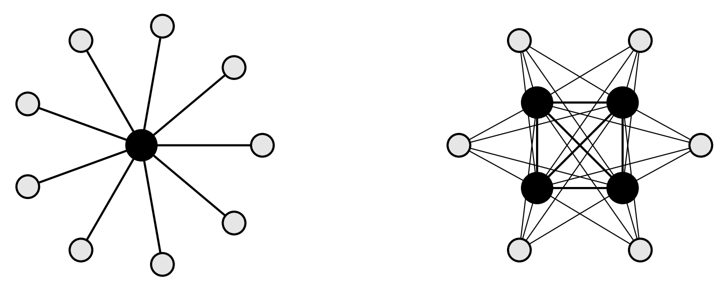

(a) star network ($k = 1$) (b) dense core ($k = 4$)

Figure 2: Comparison of CP architectures

We assume that $m > 0$ is the uniform mean influence and $v^2 \geq 0$ is the variance. We consider a central bank with a uniform solvency target $z = \mathbf{1}$ and a block-symmetric liquidity intervention strategy $\mathbf{x} = [x_c \mathbf{1}_k^\top, x_p \mathbf{1}_{n-k}^\top]^\top$, so Property A holds for the reduced mean-influence matrix. By taking into account the worst-case scenario, the DM minimizes the following objective function:

$$
\begin{aligned}
&\frac{1}{2}\mathbf{x}^\top(\mathbf{M} + \mathbf{B}^* + c\mathbf{I})\mathbf{x} - \psi^\top \mathbf{x} \\
&= \underbrace{\frac{k}{2}\Omega\left[(k-1)x_c + (n-k)x_p\right]^2}_{\text{systemic risk via core}} + \underbrace{\frac{n-k}{2}\Omega\left[kx_c\right]^2}_{\text{systemic risk via periphery}} + \underbrace{\frac{c}{2}[kx_c^2 + (n-k)x_p^2]}_{\text{intervention cost}} \\
&\quad - \underbrace{[k\psi_c x_c + (n-k)\psi_p x_p]}_{\text{linear benefit}},
\end{aligned}
$$

where $\psi_c = m(n-1)$ and $\psi_p = mk$. Differentiating with respect to x_c and x_p yields the linear system:

$$\begin{bmatrix} \Omega((n-2)k+1)+c & \Omega(k-1)(n-k) \\ \Omega k(k-1) & \Omega k(n-k)+c \end{bmatrix} \begin{bmatrix} x_c \\ x_p \end{bmatrix} = m \begin{bmatrix} n-1 \\ k \end{bmatrix}.$$

Since $\Omega, c > 0$, the coefficient matrix is invertible, yielding a unique intervention $(x_c^*, x_p^*)^\mathsf{T}$. The next proposition characterizes its key qualitative property.

Proposition 2 *The unique robust block-symmetric intervention (x_c^*, x_p^*) satisfies $x_c^* > 0$ and $x_p^* > 0$, with $x_p^* > 0$ holding even at $c = 0$. Hence the robust central bank always allocates strictly positive liquidity to the periphery, ruling out bang-bang corner solutions on the core.*

Proposition 2 provides an important insight for financial regulation. Standard network intervention models under linear budget constraints (e.g., Galeotti et al., 2020) often dictate "bang-bang" corner solutions (e.g., $x_p = 0$) where all liquidity is pumped into the highly central core. However, in our robust framework, the planner anticipates that the core also acts as the primary hub for systemic risk. At $x_p = 0$, the marginal leverage benefit toward the periphery is strictly positive and dominates the cross-coupling risk through the core at the constrained optimum x_c^*, so a robust central bank always finds it optimal to allocate at least a marginal amount of liquidity directly to the periphery ($x_p^* > 0$).

3.2.3 The Directed Bipartite Network: Supply Chain Management

We apply our framework to supply chain management, modeled as a directed bipartite network consisting of k upstream agents and $n-k$ downstream agents (see Figure 3). Upstream agents (U) represent suppliers of essential parts or raw materials, while downstream agents (D) represent manufacturers producing the end product.

Each upstream agent $j \in U$ influences every downstream agent $i \in D$ with mean $m > 0$ and variance $v^2 \geq 0$, capturing uncertainty in the transmission of inputs (e.g., due to shipping delays or yield issues). Each downstream agent $i \in D$ additionally receives a direct policy instrument with mean $\mu > 0$ and variance $v_D^2 \geq 0$, capturing direct capacity subsidies whose effectiveness varies by manufacturer. Influence flows only from U to D; no downstream agent influences upstream agents. The DM (e.g., a government planner) targets the downstream production level, setting $z_D = 1$ and

$z_U = 0$, with cost matrix $\mathbf{C} = c\mathbf{I}$ and reference $\mathbf{x}^0 = \mathbf{0}$. In addition, we focus on the regime where the intervention cost c is sufficiently large to ensure interior solutions in both blocks; this captures settings with significant budget or implementation frictions, which is often assumed in the network intervention literature (e.g., Galeotti et al., 2020; Jeong and Shin, 2024).

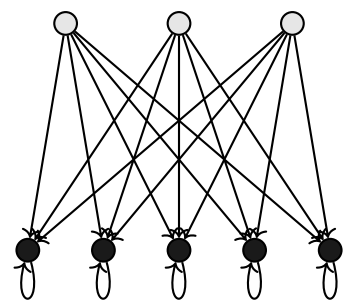

Figure 3: The directed bipartite network.

Notes: The network consists of 3 upstream agents (U) and 5 downstream agents (D). Directed arrows represent the uncertain transmission of grants from U to D with mean m and variance v^2. Self-loops at downstream agents capture direct local provision with mean μ and variance v_D^2.

By taking into account the worst-case scenario, the DM minimizes the following objective function:

$$\frac{n-k}{2}\Big[k^2\Omega\, x_U^2 + 2k\Omega_{UD}\, x_U x_D + \Omega_D\, x_D^2\Big] + \frac{c}{2}\Big[k x_U^2 + (n-k)x_D^2\Big] - (n-k)\Big[km x_U + \mu x_D\Big], \tag{10}$$

where $\Omega = m^2 + v^2$, $\Omega_D = \mu^2 + v_D^2$, and $\Omega_{UD} = m\mu + v v_D$. The first-order conditions yield the unique solution characterized as a solution of the following equation:

$$\begin{bmatrix} k^2(n-k)\Omega + kc & k(n-k)\Omega_{UD} \\ k(n-k)\Omega_{UD} & (n-k)\Omega_D + (n-k)c \end{bmatrix} \begin{bmatrix} x_U^* \\ x_D^* \end{bmatrix} = (n-k)\begin{bmatrix} km \\ \mu \end{bmatrix}. \tag{11}$$

The cross term $\Omega_{UD} = m\mu + v v_D$ captures the worst-case comovement between the upstream network channel and the direct downstream channel: adversarial Nature simultaneously amplifies uncertainty in both instruments, reducing the benefit of splitting resources across them. Consequently, Ω_{UD} depends on the *product* $v v_D$, so this adversarial coupling is active only when both channels are uncertain. This structure gives rise to two qualitatively distinct comparative-statics regimes.

Proposition 3 *For sufficiently large c, $x_U^*, x_D^* > 0$. Moreover:*

(i) Substitution regime. When $v_D = 0$, $\frac{x_U^*}{x_D^*} = \frac{(n-k)m/\mu}{1+k(n-k)v^2/c}$, *strictly decreasing in* v. *Symmetrically, when* $v = 0$, $\frac{x_U^*}{x_D^*} = \frac{(n-k)m}{\mu}\left(1 + \frac{v_D^2}{c}\right)$, *strictly increasing in* v_D.

(ii) Joint-retrenchment regime. When $v, v_D > 0$ *and* $c \geq \max\{k(n-k)m^2, \mu^2\}$, *increasing* v *near zero strictly reduces both* x_U^* *and* x_D^*.

Proposition 3 reveals a nuanced policy trade-off in supply chain management. With single-channel uncertainty (e.g., upstream shipping is volatile but downstream subsidies are administratively reliable), the robust planner substitutes toward the safer channel: increasing the uncertain channel's variance shifts allocation toward the reliable instrument, recovering the intuitive "flight to safety."

With joint uncertainty, however, this intuition breaks down. The worst-case cross-covariance $\Omega_{UD} = m\mu + vv_D$ couples the two channels: greater dispersion in either link not only raises that link's variance but also amplifies the adversary's capacity to comove shocks across both channels. The DM cannot escape this coupling by reallocating; the optimal response is *joint retrenchment*, a measured withdrawal across both instruments to limit overall exposure. Moreover, when the network multiplier dominates the direct channel ($k(n-k)m^2 \geq \mu^2$), joint uncertainty even shifts the relative allocation *toward* the high-leverage network instrument, as direct subsidies retreat more rapidly than network subsidies.

The implication is sharp: in a moderate disruption affecting only one channel (e.g., a localized supply bottleneck with reliable government intervention capacity), the planner can substitute reliably toward the unaffected channel; in a deep systemic crisis where multiple channels are simultaneously exposed (e.g., global shipping bottlenecks coinciding with local administrative disruptions), the planner must retrench across all instruments because the adversary can coordinate shocks across them. Robust policy in dual-uncertainty environments reduces exposure rather than reallocating across channels.

3.2.4 Network Expansion: Recruiting New Faculty

In the context of faculty recruitment, the relevant network evolves as a new researcher joins the institution. The DM (a university administrator allocating research support) continues to evaluate interventions using the max–min objective (1), but the information available about within-row covariances expands only partially.

From a fixed network to an expanding network. The baseline analysis treats the set of agents N as fixed and models uncertainty through within-row dependence ambiguity. In many applications the relevant network evolves: new agents, locations, or institutions enter, and the DM must update the intervention while retaining only partial information about how the new links comove with existing ones. This subsection studies network expansion within the same robust framework. The DM continues to evaluate interventions using the max–min objective (1) and row-wise uncertainty, but the information available about within-row covariances expands only partially when a new node arrives. We show that the tractability of the baseline model carries over and that worst-case dependence can still be characterized in closed form given the available moments.

Setup. Maintain the objective (1) and the row-wise ambiguity structure from Subsection 2.2. The agent set expands from $N = \{1, \dots, n\}$ to $N' = N \cup \{n+1\}$. The DM observes the new link-wise moments involving the entrant but has only partial information about within-row covariances that involve $n+1$. Formally, for each receiver $i \in N'$, Nature chooses a covariance matrix over the incoming shocks to i subject to (i) positive semi-definiteness and (ii) the moment information available to the DM. The benchmark fixed-network case corresponds to knowing only the diagonal (variances) and leaving all within-row covariances free; network expansion introduces additional linear restrictions for receivers $i \in N$ coming from the DM's prior information about the pre-expansion network.

The key difference from Subsection 3.1 is therefore informational: for some receivers $i \in N$, the uncertainty set for the extended covariance matrix is a restricted subset of the elliptope determined by the known principal submatrix for the pre-expansion agents. The new receiver $n+1$ remains in the baseline case with unrestricted within-row dependence, so Nature's worst-case choice for $i = n+1$ continues to exhibit the rank-1 structure from Theorem 1.

To model the DM's partial information, we introduce the following notation. Let $\mathbf{B}_i \in \mathrm{PSD}^n$ for $1 \le i \le n$ represent the covariance matrix among the existing n agents, and let $\overline{\mathbf{B}}_i \in \mathrm{PSD}^{n+1}$ for $1 \le i \le n+1$ denote the covariance matrix of all $n+1$ agents. Since we are modeling the DM's partial information, we require that for each agent $i = 1, \dots, n+1$, $\mathbf{B}_i$ is the $n \times n$ principal submatrix of $\overline{\mathbf{B}}_i$, obtained by selecting the first n rows and columns of $\overline{\mathbf{B}}_i$. On the other hand, let $b_i = \left(\overline{\mathbf{B}}_i\right)_{(n+1)\,(n+1)} > 0$

represent the variance of the link $\mathbf{G}_{i(n+1)}$ for $i = 1, \ldots, n+1$. Given $\mathbf{B} \in \mathrm{PSD}^n$ and $b > 0$, we define $\overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B},b}$ as the set of all extended matrices of $\mathbf{B}$ with its last $(n+1, n+1)$ entry fixed as b, forming a new uncertainty set. Specifically, we define $\overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B},b} = \overline{\mathcal{B}}_{\mathbf{B},b} \cap \mathrm{PSD}^{n+1}$, where $\overline{\mathcal{B}}_{\mathbf{B},b}$ is defined as

$$\overline{\mathcal{B}}_{\mathbf{B},b} = \{\overline{\mathbf{B}} \in \mathbb{R}^{(n+1)\times(n+1)} \mid \overline{\mathbf{B}} \text{ is symmetric, } \overline{\mathbf{B}}_{ij} = \mathbf{B}_{ij} \text{ for all } 1 \leq i, j \leq n, \overline{\mathbf{B}}_{(n+1)(n+1)} = b\}.$$

We denote by $\overline{\mathbf{x}} \in \mathbb{R}^{n+1}$ the intervention of the DM.

Row-wise uncertainty sets under expansion. The DM's pre-expansion information pins down the $n \times n$ principal submatrix of $\overline{\mathbf{B}}_i$ to be $\mathbf{B}_i$, while the covariances involving the entrant are left free subject to positive semi-definiteness and the known variance b_i of the entrant link. Accordingly, Nature's feasible set for receiver $i \in N$ is $\overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B}_i,b_i}$.

For the entrant receiver $n+1$, we maintain the baseline moment-information structure: the DM knows the marginal variances of the incoming links but has no additional information about within-row covariances. Let $\mathcal{B}_{n+1}$ denote the corresponding uncertainty set defined as in Subsection 2.2 (with dimension $n+1$).

Given an intervention $\overline{\mathbf{x}} \in \mathbb{R}^{n+1}$, Nature's inner problem under expansion decomposes across receivers as in the baseline model:

$$\max_{\{\overline{\mathbf{B}}_i\}} \sum_{i=1}^{n+1} \frac{1}{2} \langle \overline{\mathbf{x}}, \overline{\mathbf{B}}_i \overline{\mathbf{x}} \rangle \quad \text{s.t.} \quad \overline{\mathbf{B}}_i \in \overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B}_i,b_i} \text{ for } i \in N, \quad \overline{\mathbf{B}}_{n+1} \in \mathcal{B}_{n+1}.$$

Proposition 4 establishes the uniqueness of Nature's best response with respect to the DM's intervention $\overline{\mathbf{x}}$, that is, the uniqueness of the worst-case scenario.[13]

Proposition 4 *Fix $i \in N$ and suppose $\mathbf{B}_i \in \mathrm{PD}^n$. If $\overline{\mathbf{x}}$ has no zero entry, then the maximizer of $\langle \overline{\mathbf{x}}, \overline{\mathbf{B}}_i \overline{\mathbf{x}} \rangle$ over $\overline{\mathbf{B}}_i \in \overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B}_i,b_i}$ is unique.*[14]

Interpretation for faculty recruitment: Under network expansion, the new hire $n+1$ represents a source of unconstrained, rank-1 ambiguity. Nature can perfectly align the new agent's uncertain collaborative shocks to inflict maximum variance.

[13]For expositional simplicity, we assume here that the DM possesses no knowledge of the new agent $n+1$'s interaction with the existing n agents. As such, Proposition 4 is a special case of Proposition A3, in which the DM is allowed to have "partial" knowledge of the new agent $n+1$'s interaction with the existing n agents. See Appendix C for more details.

[14]To build intuition for this result, Online Appendix OA.1 provides a geometric illustration of a two-agent network expanding to include a single entrant. We demonstrate how the DM's partial correlation constraints generate strictly convex uncertainty sets that uniquely determine Nature's worst-case response.

Conversely, Nature is "handcuffed" by the known collaboration history among existing faculty. Consequently, a robust administrator facing high structural ambiguity will naturally tilt research funding $\overline{\mathbf{x}}$ heavily toward the established members, acting cautiously toward the new hire until their collaborative covariances are revealed through actual interaction.

4 Comparative Statics

4.1 The Value of Robust Intervention

We now define the value of robust intervention. As a benchmark, we consider the standard approach in the network intervention literature (e.g., Galeotti et al., 2020), which optimizes against the expected network structure. We refer to this as the *naive intervention*, denoted $\mathbf{x}^N$. This DM relies solely on the first moments of the network, minimizing the objective function while treating the network as if it were deterministic at its mean. Formally, $\mathbf{x}^N$ solves:

$$\mathbf{x}^N = \arg\min_{\mathbf{x}\in\mathbb{R}^n} \frac{1}{2}\left(\langle \mathbf{x}, \mathbf{M}\mathbf{x}\rangle + \langle \mathbf{x}, \mathbf{C}\mathbf{x}\rangle\right) - \langle \psi^0 + \psi, \mathbf{x}\rangle.$$

This corresponds to setting $\mathbf{v}_{ij} = 0$ for all $i, j \in N$ in the robust problem, so that $\mathbf{B} = \mathbf{0}$ and all variance is ignored entirely. The unique solution is:

$$\mathbf{x}^N = (\mathbf{M} + \mathbf{C})^{-1}(\psi^0 + \psi). \tag{12}$$

The robust intervention, by contrast, accounts for the worst-case dependence structure and satisfies $\mathbf{x}^* = (\mathbf{M}+\mathbf{B}^*+\mathbf{C})^{-1}(\psi^0+\psi)$. Since $\mathbf{B}^* \succeq 0$, we have $\mathbf{M}+\mathbf{B}^*+\mathbf{C} \succeq \mathbf{M}+\mathbf{C}$, which implies $(\psi^0 + \psi)^\top(\mathbf{x}^N - \mathbf{x}^*) \geq 0$, so the robust DM intervenes less aggressively than the naive DM in the direction of $\psi^0 + \psi$.

Given $\mathbf{x}^N$, Nature's worst-case response to the naive intervention is $\mathbf{B}^N = \sum_{i=1}^n \mathbf{q}_i(\mathbf{x}^N) \otimes \mathbf{q}_i(\mathbf{x}^N)$, where $\mathbf{q}_i(\mathbf{x}^N) = (s(x_j^N)\mathbf{v}_{ij})_{j=1}^n$ as in the main theorem. Again, each row-wise component is rank-1 and aligned with the sign pattern of $\mathbf{x}^N$, so $\mathbf{B}^N$ is Nature's adversarial row-wise rank-1 response to the naive DM. We define two measures of the value of robust intervention.

The *ex-ante value* of robust intervention compares the minimax guarantees, each evaluated against Nature's worst-case response:

$$\mathcal{V} = f(\mathbf{x}^N, \mathbf{B}^N) - f(\mathbf{x}^*, \mathbf{B}^*). \tag{13}$$

This value also represents the *ex-ante cost of uncertainty*: it measures the total loss suffered by the naive DM relative to the robust DM when Nature is allowed to fully exploit the lack of precaution. Since $\mathbf{x}^*$ minimizes $\max_{\mathbf{B}\in\mathcal{B}} f(\mathbf{x},\mathbf{B})$ by construction, $\mathcal{V}\geq 0$.

The *ex-post value* of robust intervention measures the realized benefit of the robust strategy relative to the naive one, given the worst-case network $\mathbf{B}^*$ that the robust DM guarded against:

$$\mathcal{V}' = f(\mathbf{x}^N,\mathbf{B}^*) - f(\mathbf{x}^*,\mathbf{B}^*). \tag{14}$$

Similarly, $\mathcal{V}'$ characterizes the *ex-post cost (or regret) of uncertainty*: the loss from ignoring uncertainty, conditional on the worst-case network $\mathbf{B}^*$ being realized. It measures how much worse the naive allocation $\mathbf{x}^N$ performs than the robust allocation $\mathbf{x}^*$ on this network. Since $\mathbf{x}^*$ minimizes $f(\cdot,\mathbf{B}^*)$ by its first-order condition, $\mathcal{V}'\geq 0$.

The following proposition characterizes and relates both measures.

Proposition 5 *Assume Properties A and B. Then, the values are*

$$\mathcal{V}' = \frac{1}{2}\langle \mathbf{x}^N-\mathbf{x}^*,\, (\mathbf{M}+\mathbf{B}^*+\mathbf{C})(\mathbf{x}^N-\mathbf{x}^*)\rangle \geq 0, \tag{15}$$

$$\mathcal{V} = \frac{1}{2}\langle \mathbf{x}^N,\, \mathbf{B}^N\mathbf{x}^N\rangle - \frac{1}{2}\langle \mathbf{x}^N,\mathbf{B}^*\mathbf{x}^*\rangle \geq \mathcal{V}'. \tag{16}$$

Both equal zero if and only if $\mathbf{v}_{ij}=0$ for all $i,j\in N$. Moreover, if $\mathbf{x}^N$ and $\mathbf{x}^*$ lie in the same orthant so that $\mathbf{B}^N=\mathbf{B}^*=\mathbf{B}$, then $\mathcal{V}=\frac{1}{2}\langle \mathbf{x}^N,\, \mathbf{B}(\mathbf{x}^N-\mathbf{x}^*)\rangle$.

The ordering $\mathcal{V}\geq\mathcal{V}'\geq 0$ reflects the difference in what Nature is allowed to do. In $\mathcal{V}$, Nature re-optimizes against $\mathbf{x}^N$, choosing $\mathbf{B}^N$ to maximally exploit the naive DM's strategy. In $\mathcal{V}'$, Nature is held fixed at $\mathbf{B}^*$, which was designed to exploit $\mathbf{x}^*$, not $\mathbf{x}^N$. The gap $\mathcal{V}-\mathcal{V}'=\frac{1}{2}\langle \mathbf{x}^N,(\mathbf{B}^N-\mathbf{B}^*)\mathbf{x}^N\rangle\geq 0$ is Nature's *strategic re-optimization gain* against the naive DM: it captures the additional harm Nature can inflict by adjusting its worst-case scenario from $\mathbf{B}^*$ to $\mathbf{B}^N$.

The ex-post regret $\mathcal{V}'$ has an economic interpretation: it is the weighted squared distance between the naive and robust interventions, with weight matrix $\mathbf{M}+\mathbf{B}^*+\mathbf{C}$. Since $(\mathbf{M}+\mathbf{C})(\mathbf{x}^N-\mathbf{x}^*)=\mathbf{B}^*\mathbf{x}^*$ from the two first-order conditions, the distance $\mathbf{x}^N-\mathbf{x}^*$ is entirely driven by the worst-case uncertainty $\mathbf{B}^*\mathbf{x}^*$. Thus $\mathcal{V}'$ becomes large when the robust DM makes a substantially different allocation from the naive DM, which occurs when uncertainty is high and the rank-1 correction $\mathbf{B}^*$ is large.

We now analyze how the benefits of robustness evolve as network uncertainty increases. We define the *marginal value of robust intervention* as the increase in the value of the robust strategy resulting from a marginal increase in the variance $\mathbf{v}_{ij}$

of link (i, j). Because this marginal value represents the performance loss avoided by the robust DM relative to the naive one, it also characterizes the *marginal cost of uncertainty*. Formally, the ex-ante and ex-post marginal values are given by $\frac{\partial \mathcal{V}}{\partial \mathbf{v}_{ij}}$ and $\frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}}$, respectively. The following proposition characterizes both measures, establishing their general relationship and their common value under a regularity condition.

Proposition 6 *Assume Properties A and B. Then, the ex-ante and ex-post marginal values of robust intervention are given by*

$$\frac{\partial \mathcal{V}}{\partial \mathbf{v}_{ij}} = \Big(\sum_{l \in N} \mathbf{v}_{il} |x_l^N|\Big) |x_j^N| - \Big(\sum_{l \in N} \mathbf{v}_{il} |x_l^*|\Big) |x_j^*|, \tag{17}$$

$$\frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}} = \Big(\sum_{l \in N} s(x_l^*) \mathbf{v}_{il} x_l^N\Big) s(x_j^*) x_j^N - \Big(\sum_{l \in N} \mathbf{v}_{il} |x_l^*|\Big) |x_j^*|. \tag{18}$$

The ex-ante marginal value is weakly greater than the ex-post marginal value as

$$\frac{\partial \mathcal{V}}{\partial \mathbf{v}_{ij}} - \frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}} = 2|x_j^N| \sum_{\substack{l \in N \\ s(x_l^N)s(x_l^*) \neq s(x_j^N)s(x_j^*)}} \mathbf{v}_{il} |x_l^N| \geq 0. \tag{19}$$

Moreover, if $\mathbf{x}^N$ and $\mathbf{x}^*$ lie in the same orthant, then $\mathbf{B}^N = \mathbf{B}^*$ and the two marginal values coincide as

$$\frac{\partial \mathcal{V}}{\partial \mathbf{v}_{ij}} = \frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}} = \Big(\sum_{l \in N} \mathbf{v}_{il} |x_l^N|\Big) |x_j^N| - \Big(\sum_{l \in N} \mathbf{v}_{il} |x_l^*|\Big) |x_j^*|. \tag{20}$$

Several implications follow from Proposition 6. First, the ex-ante marginal value of robust intervention ($\frac{\partial \mathcal{V}}{\partial \mathbf{v}_{ij}}$) depends only on the absolute intervention levels $|x_l^N|$ and $|x_l^*|$, not on their signs. This is because $\mathcal{V}$ compares each DM against their own worst-case Nature, which naturally aligns with the respective sign pattern of their intervention. In this sense, the *marginal ex-ante cost of uncertainty* is strictly determined by the absolute magnitude of the naive DM's over-exposure. Since it maximizes mean influence, the naive DM allocates excessive resources to highly connected network channels (i.e., $|x_l^N|$). The robust DM accounts for the fact that Nature concentrates correlation risk on these same channels and scales back the corresponding allocations (i.e., $|x_l^*|$).

In contrast, the ex-post marginal value ($\frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}}$) depends on $s(x_l^*)x_l^N$, which is the naive intervention projected onto the sign pattern of the robust intervention. When $\mathbf{x}^N$ and $\mathbf{x}^*$ share the same sign for component l, this term equals $|x_l^N|$; however, when their directions differ, it equals $-|x_l^N|$, which strictly reduces the marginal value. Consequently, $\frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}}$ can be negative. Economically, this means the *marginal ex-post regret* (or the realized marginal cost) may actually decrease as uncertainty $\mathbf{v}_{ij}$ rises. This occurs because the realized worst-case network $\mathbf{B}^*$ is fundamentally less adversarial against the naive DM when $\mathbf{x}^N$ and $\mathbf{x}^*$ point in different directions.

Second, the gap $\frac{\partial(\mathcal{V}-\mathcal{V}')}{\partial \mathbf{v}_{ij}} \geq 0$ captures Nature's additional strategic flexibility when the two interventions do not perfectly align in direction. This gap equals zero whenever $\mathbf{x}^N$ and $\mathbf{x}^*$ share the same sign pattern across all components l (which holds, in particular, when they lie in the same orthant). However, when the directional alignment $\{s(x_l^N)s(x_l^*)\}_{l\in N}$ is non-uniform, the gap is strictly positive. Increasing $\mathbf{v}_{ij}$ widens the difference between the ex-ante and ex-post marginal values of robust intervention, because Nature's ability to re-optimize from $\mathbf{B}^*$ to $\mathbf{B}^N$ is amplified by components where the naive and robust DMs disagree on the direction of intervention.

Third, when $\mathbf{x}^N$ and $\mathbf{x}^*$ lie in the same orthant, the ex-ante and ex-post marginal values of robust intervention coincide and admit a transparent decomposition: the term $(\sum_l \mathbf{v}_{il}|x_l^N|)|x_j^N|$ represents the marginal worst-case exposure of the naive DM along link (i,j), while $(\sum_l \mathbf{v}_{il}|x_l^*|)|x_j^*|$ represents the corresponding exposure of the robust DM. The difference between these two exposures perfectly isolates the marginal cost of uncertainty, demonstrating that it is driven entirely by the naive DM's relative over-exposure. This same-orthant condition arises naturally in many economic environments, including all four network applications analyzed in Subsection 3.2, where the interventions $\mathbf{x}^N$ and $\mathbf{x}^*$ are strictly positive.

4.2 Applications: Revisited

d-regular network. In the d-regular network, the symmetric structure yields closed-form expressions for both the naive and robust interventions, $x^N = \frac{dm}{d^2m^2+c} > x^*(d) = \frac{dm}{d^2\Omega+c}$, where $\Omega = m^2 + v^2 \geq m^2$. Because $m > 0$, both interventions are strictly positive and lie in the same orthant, so $\mathbf{B}^N = \mathbf{B}^*$ by Proposition 6. Furthermore, $|x_l^N| \geq |x_l^*|$ holds component-wise. Consequently, $f(\mathbf{x}^N, \mathbf{B}^N) = f(\mathbf{x}^N, \mathbf{B}^*)$ and $\mathcal{V} = \mathcal{V}'$: the ex-ante and ex-post values of robust intervention coincide, and Nature has no additional strategic flexibility against the naive planner. As such, the ex-ante and ex-post marginal values for any link (i,j) perfectly coincide, isolating the marginal cost of uncertainty as the precise gap between naive over-exposure and robust mitigation:

$$\frac{\partial \mathcal{V}}{\partial \mathbf{v}_{ij}} = \frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}} = dv\left((x^N)^2 - (x^*)^2\right) > 0.$$

Figure 4 plots the two worst-case objectives as functions of the uncertainty level v. Three features are notable. First, at $v = 0$, the objectives coincide and $\mathcal{V} = \mathcal{V}' = 0$: without uncertainty, the naive and robust planners are identical. Second, $f(\mathbf{x}^*, \mathbf{B}^*)$ increases monotonically in v: the robust planner's guaranteed performance deteriorates as uncertainty grows. Third, $f(\mathbf{x}^N, \mathbf{B}^N)$ increases much more steeply in v and crosses zero at $v = \sqrt{m^2 + c/d^2}$: beyond this threshold, the naive planner's worst-case scenario turns into outright value destruction. The gap between the two objectives ($\mathcal{V}$) therefore widens rapidly, demonstrating that the ex-ante value of robust intervention—and consequently the economic cost of uncertainty—is most pronounced in highly volatile environments.

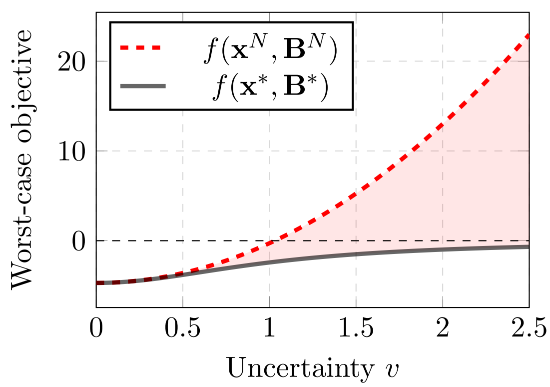

Figure 4: Worst-case objective values for the d-regular network.

Notes: $m = 1$, $c = 1$, $n = 10$, $d = 4$. The shaded region represents $\mathcal{V} = \mathcal{V}'$, and the naive planner's worst-case performance crosses into value destruction at $v = \sqrt{m^2 + c/d^2} \approx 1.03$.

CP network. In the core-periphery network $CP_{n,k}$, we distinguish two types of link uncertainty: v_1 for core-to-core (C-C) links and v_2 for core-to-periphery (C-P) links. Figure 5 plots the worst-case objectives $f(\mathbf{x}^*, \mathbf{B}^*)$ and $f(\mathbf{x}^N, \mathbf{B}^N)$ as functions of v_1 (left panel, fixing $v_2 = 0$) and v_2 (right panel, fixing $v_1 = 0$). Similar to the d-regular case, both interventions lie in the same orthant and $|x_l^N| \geq |x_l^*|$ holds componentwise. This ensures that the ex-ante and ex-post values of robust intervention perfectly coincide ($\mathcal{V} = \mathcal{V}'$), allowing the gap between the two objectives to visually capture the exact cost of uncertainty, driven entirely by the penalty of naive over-exposure.

The two panels reveal a fundamental asymmetry between C-C and C-P uncertainty: increasing v_2 generates a substantially larger gap ($\mathcal{V}$) than increasing v_1 by the same amount. To see why, recall from Proposition 6 that the marginal value of robust intervention for link (i, j) is $(\sum_l \mathbf{v}_{il}|x_l^N|)|x_j^N| - (\sum_l \mathbf{v}_{il}|x_l^*|)|x_j^*|$, which depends

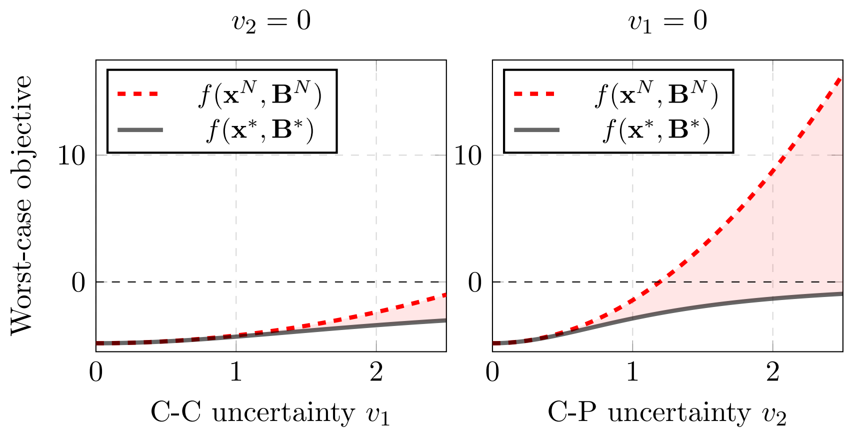

Figure 5: Worst-case objective values for the core-periphery network.

Notes: $CP_{10,3}$ ($m = 1$, $c = 1$). The shaded regions represent $\mathcal{V} = \mathcal{V}'$. C-P uncertainty generates a substantially larger cost of uncertainty than C-C uncertainty for the same variance level.

strictly on the row-i exposure $\sum_l \mathbf{v}_{il}|x_l|$. A core agent's row aggregates over both blocks—$v_1(k-1)|x_c| + v_2(n-k)|x_p|$—while a periphery agent's row aggregates over the core only—$v_2 k|x_c|$. Hence, v_1 enters only the core row's exposure and governs the $k(k-1)$ C-C links, whereas v_2 enters both row types and governs the $2k(n-k)$ C-P links. Because each C-P link spans the boundary between the two blocks, its variance propagates Nature's uncertainty through both core and periphery receivers simultaneously, while a C-C link's variance is confined to the core block. Consequently, cross-block links disproportionately amplify the naive planner's over-exposure, driving a much steeper marginal cost of uncertainty.

This asymmetry reflects a general principle embedded in Proposition 6: links that connect agents across distinct blocks generate a strictly larger marginal value of robust intervention, because Nature can simultaneously correlate their variance across a broader set of receivers. In the core-periphery architecture, C-P links act as the systemic conduit between the two segments, making their uncertainty the dominant source of the overall value of robustness. This finding yields a direct policy implication: a regulator facing limited capacity to measure network uncertainty should prioritize estimating the variance of cross-block links over within-block links, as the former contribute disproportionately to the naive planner's over-exposure and the resulting economic cost of uncertainty.

Directed bipartite network. In the directed bipartite network, we distinguish two sources of uncertainty: v for upstream-to-downstream links and v_D for the downstream self-loops. The naive intervention $\mathbf{x}^N = (x_U^N, x_D^N)^\intercal$ is obtained by setting $v = v_D = 0$ in system (11). Since $m, \mu > 0$, both interventions lie in the same orthant, ensuring that the ex-ante and ex-post values of robust intervention perfectly coincide ($\mathcal{V} = \mathcal{V}'$). Unlike the uniform shrinking seen in the previous symmetric models, the robust planner in a bipartite network may substitute allocations across channels rather than uniformly reducing them. However, the robust planner guarantees that the overall row-wise exposure $\sum_l \mathbf{v}_{il}|x_l^*|$ remains strictly smaller under $\mathbf{x}^*$. As a result, the ex-ante and ex-post marginal values coincide and remain strictly positive, capturing the exact marginal cost of the naive planner's over-exposure. With respect to v and v_D, these marginal values are:[15]

$$\frac{\partial \mathcal{V}}{\partial v} = (n-k)k^2 v\Big[(x_U^N)^2 - (x_U^*)^2\Big] + (n-k)k\, v_D\Big[x_U^N x_D^N - x_U^* x_D^*\Big], \tag{21}$$

$$\frac{\partial \mathcal{V}}{\partial v_D} = (n-k)k\, v\Big[x_U^N x_D^N - x_U^* x_D^*\Big] + (n-k)v_D\Big[(x_D^N)^2 - (x_D^*)^2\Big]. \tag{22}$$

These marginal values admit a transparent decomposition into two distinct channels: (i) a *direct* channel that scales the squared intervention gap for the agents whose links carry the specific variance, and (ii) a *cross* channel that captures the covariance between the upstream and downstream intervention gaps, weighted by the complementary variance parameter. When $v_D = 0$ or $v = 0$, the cross channel vanishes and the marginal value reduces strictly to the direct channel. In these regimes, the equations cleanly isolate the independent marginal cost of uncertainty, demonstrating how the naive planner's over-exposure is penalized within a single segment of the supply chain before cross-channel effects compound the damage.

The ratio of the direct marginal values at equal variance $v = v_D$ is $k^2 \frac{(x_U^N)^2 - (x_U^*)^2}{(x_D^N)^2 - (x_D^*)^2} > 1$, reflecting two structural forces: k^2 upstream links per downstream agent versus only one self-loop, and a fundamentally larger naive over-exposure in the upstream direction when the network multiplier km is large. Because both forces scale with k, upstream link variance disproportionately amplifies the marginal cost of uncertainty. Consequently, a regulator facing limited measurement capacity should strictly

[15] By Proposition 6, $\frac{\partial \mathcal{V}}{\partial v} = \sum_{d \in D} \sum_{u \in U} \frac{\partial \mathcal{V}}{\partial \mathbf{v}_{du}}$. For each downstream agent d, the row-d exposure is $\sum_l \mathbf{v}_{dl}|x_l| = kv \cdot |x_U| + v_D \cdot |x_D|$, so

$$\frac{\partial \mathcal{V}}{\partial \mathbf{v}_{du}} = (kv\, x_U^N + v_D\, x_D^N)\, x_U^N - (kv\, x_U^* + v_D\, x_D^*)\, x_U^*,$$

using same-orthant positivity. Summing over $(n-k) \times k$ such links yields the expression of $\frac{\partial \mathcal{V}}{\partial v}$. The derivation of $\frac{\partial \mathcal{V}}{\partial v_D}$ is analogous, with the sum running over the $n-k$ self-loops.

prioritize estimating upstream variance over downstream self-loops.

Figure 6 illustrates these two regimes ($n = 8$, $k = 3$, $m = \mu = 1$, $c = 1$). In both panels, the worst-case objectives deteriorate as uncertainty increases. Specifically, $f(\mathbf{x}^N, \mathbf{B}^*)$ climbs toward zero and beyond, while $f(\mathbf{x}^*, \mathbf{B}^*)$ increases much more slowly, causing the values of robust intervention ($\mathcal{V} = \mathcal{V}'$) to steadily widen. The critical contrast is the rate: in the left panel (upstream uncertainty), $f(\mathbf{x}^N, \mathbf{B}^*)$ increases steeply and crosses zero, resulting in outright value destruction for the naive planner. In the right panel (downstream uncertainty), the objective is nearly flat, visually demonstrating that the economic cost of uncertainty in this architecture is overwhelmingly driven by the upstream supply channels.

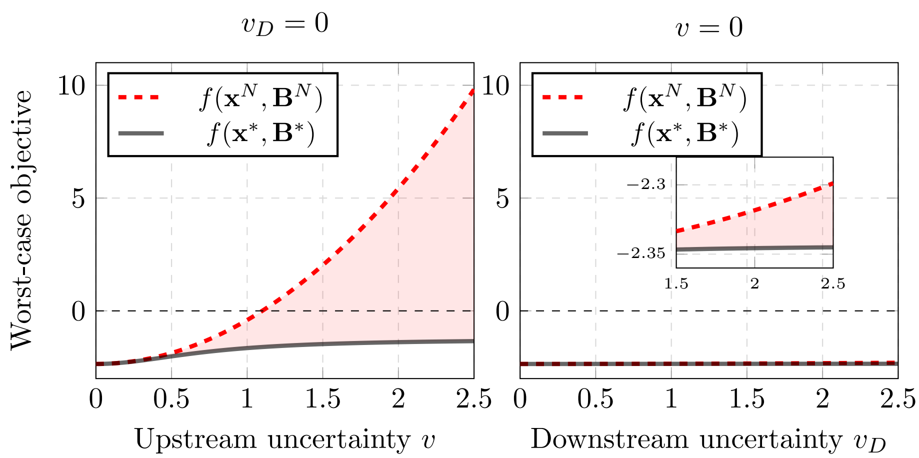

Figure 6: Worst-case objective values for the directed bipartite network.

Notes: $n = 8$, $k = 3$, $m = \mu = 1$, $c = 1$. The shaded regions represent $\mathcal{V} = \mathcal{V}'$. Upstream uncertainty generates a substantially larger cost than downstream uncertainty. The inset magnifies $v_D \in [1.5, 2.5]$ to show that a small gap still emerges downstream.

5 Conclusion

This paper develops a framework for robust intervention in networks where the decision maker faces structural uncertainty about how agents interact. By modeling this environment as a zero-sum game between the DM and an adversarial Nature, we show that worst-case correlation risk exhibits a row-wise rank-1 structure. This reduces the ambiguity over the network's joint distribution to a fixed-point problem driven by the sign pattern of the intervention.

Across varied economic environments, our findings collectively highlight a fundamental shift required for optimal network policy. Standard deterministic models advocate targeting high-influence agents to maximize systemic leverage. However, we demonstrate that this naive approach creates a severe over-exposure to correlation risk. Whether a health authority is managing pandemic contacts, a central bank is injecting liquidity, or a government is subsidizing supply chains, a robust planner must often temper uniform interventions and actively support peripheral agents to hedge against the vulnerabilities introduced by highly connected nodes. Furthermore, we show that the resulting economic cost of uncertainty is disproportionately driven by cross-block and upstream links, which act as systemic conduits for Nature's worst-case dependence.

Ultimately, effective network interventions must explicitly balance the mean influence of an action against the correlation risk it exposes the system to. Future work could expand this framework to dynamic settings where network structures evolve endogenously over time. Additionally, extending our approach to the graphon framework (Parise and Ozdaglar, 2023) would allow for the robust analysis of continuous, large-scale networks, offering a powerful tool for policy design under deep structural uncertainty.

References

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608.

Allouch, N. (2015). On the private provision of public goods on networks. *Journal of Economic Theory*, 157:527–552.

Belhaj, M., Deroïan, F., and Safi, S. (2023). Targeting in networks under costly agreements. *Games and Economic Behavior*, 140:154–172.

Ben-Tal, A., Ghaoui, L. E., and Nemirovski, A. (2009). *Robust Optimization*. Princeton University Press.

Bergemann, D. and Morris, S. (2005). Robust mechanism design. *Econometrica*, 73(6):1771–1813.

Borgatti, S. P. and Everett, M. G. (1999). Models of core/periphery structures. *Social Networks*, 21(4):375–395.

Bramoullé, Y. and Kranton, R. (2007). Public goods in networks. *Journal of Economic Theory*, 135(1):478–494.

Brooks, B. and Du, S. (2021). Maxmin auction design with known expected values. *Working Paper*.

Brooks, B. and Du, S. (2024). On the structure of informationally robust optimal mechanisms. *Econometrica*, 92(5):1391–1438.

Carroll, G. (2015). Robustness and linear contracts. *American Economic Review*, 105:536–563.

Chaudhuri, P. K., Jackson, M. O., Sarangi, S., and Tzavellas, H. (2024). Games under network uncertainty. *Working Paper*.

Che, E. (2022). Distributionally robust optimal auction design under mean constraints.

Craig, B. and von Peter, G. (2014). Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23(3):322–347.

Cremer, J. and McLean, R. P. (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257.

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

Delage, E. and Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612.

Della Lena, S. (2024). The spread of misinformation in networks with individual and social learning. *European Economic Review*, 168:104804.

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

Frankel, A. (2014). Aligned delegation. *Review of Economic Studies*, 104:66–83.

Fricke, D. and Lux, T. (2015). Core-periphery structure in the overnight money market: Evidence from the e-mid trading platform. *Computational Economics*, 45(3):359–395.

Galeotti, A., Golub, B., and Goyal, S. (2020). Targeting Interventions in Networks. *Econometrica*, 88(6):2445–2471.

Galeotti, A., Golub, B., Goyal, S., Talamás, E., and Tamuz, O. (2024). Robust Market Interventions. *Working Paper*.

Galeotti, A. and Goyal, S. (2009). Influencing the influencers: A theory of strategic diffusion. *Rand Journal of Economics*, 40(3):509–532.

Galeotti, A., Goyal, S., Jackson, M. O., Vega-Redondo, F., and Yariv, L. (2010). Network games. *Review of Economic Studies*, 77(1):218–244.

Garrett, D. F. (2014). Robustness of simple menus of contracts in cost-based procurement. *Games and Economic Behavior*, 87:631–641.

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

Golub, B. and Jackson, M. O. (2012). How homophily affects the speed of learning and best-response dynamics. *Quarterly Journal of Economics*, 127(3):1287–1338.

Guo, Y. and Shmaya, E. (2025). Robust monopoly regulation. *American Economic Review*, 115(2):599–634.

Gupta, V. and Kallus, N. (2022). Data pooling in stochastic optimization. *Management Science*, 68(3):1595–1615.

He, W. and Li, J. (2022). Correlation-robust auction design. *Journal of Economic Theory*, 200:105403.

Jackson, M. O. (2010). *Social and Economic Networks*. Princeton University Press.

Jeong, D. and Shin, E. (2024). Optimal Influence Design in Networks. *Journal of Economic Theory*, 220:105877.

Kambhampati, A. (2024). Robust performance evaluation of independent agents. *Theoretical Economics*, 19:1151–1184.

Koçyiğit, Ç., Iyengar, G., Kuhn, D., and Wiesemann, W. (2019). Distributionally robust mechanism design. *Management Science*, 66(1):159–189.

Li, W. and Tan, X. (2025). Dynamic Network Influence: The Art of Strategic Messaging. *Working Paper*.

Liang, N. (2018). Eight lessons for fighting the next financial crisis. *The Brookings Institution*. Accessed: 2024-12-30.

Meyer, C. D. (2010). *Matrix Analysis and Applied Linear Algebra*. SIAM: Society for Industrial and Applied Mathematics.

Montiel Olea, J. L., Qiu, C., and Stoye, J. (2025). Decision theory for treatment choice problems with partial identification. *Review of Economic Studies*. Forthcoming. arXiv:2312.17623.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–71.

Parise, F. and Ozdaglar, A. (2023). Graphon games: A statistical framework for network games and interventions. *Econometrica*, 91(1):191–225.

Shin, E. (2021). Social network formation and strategic interaction in large networks. *Mathematical Social Sciences*, 111:34–54.

Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics*, 166(1):138–156.

Sun, Y., Zhao, W., and Zhou, J. (2023). Structural intervention in networks. *International Economic Review*, 64(4):1533–1563.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442.

Zymler, S., Kuhn, D., and Rustem, B. (2013). Distributionally robust joint chance constraints with second-order moment information. *Mathematical Programming*, 137(1):167–198.

A Decomposition

We here provide the decomposition of the DM's objective in (1) into mean, uncertainty, and target-alignment components, and collects omitted algebraic derivations used in the main text. Let $\mathrm{E}_\pi[\mathbf{G}_{ij}] = \mathbf{m}_{ij}$ and $\mathrm{Var}_\pi(\mathbf{G}_{ij}) = \mathbf{v}_{ij}^2$ for all $i, j \in N$. For each pair of agents i and j, let $\mathbf{U}_{ij} = \mathbf{G}_{ij} - \mathbf{m}_{ij}$ denote the deviation of the influence from its mean. By construction, $\mathrm{E}_\pi[\mathbf{U}_{ij}] = 0$ and $\mathrm{Var}_\pi(\mathbf{U}_{ij}) = \mathbf{v}_{ij}^2$ for all $i, j \in N$. Let $\mathbf{m}_i = \mathrm{E}_\pi[\mathbf{G}_i^\top] \in \mathbb{R}^n$ be the vector of mean influences toward agent i. Define $\mathbf{U}_i = \mathbf{G}_i^\top - \mathbf{m}_i \in \mathbb{R}^n$ as the vector of deviations of these influences. Then the expected squared deviation between the outcome for agent i and its target z_i is

$$\begin{aligned}\mathrm{E}_\pi\left[|\mathbf{G}_i\mathbf{x} - z_i|^2\right] &= \sum_{j=1}^n \sum_{k=1}^n x_j\big(\mathbf{m}_{ij}\mathbf{m}_{ik} + \mathrm{E}_\pi[\mathbf{U}_{ij}\mathbf{U}_{ik}]\big)x_k - 2z_i \sum_{l=1}^n \mathbf{m}_{il}x_l + z_i^2 \\ &= \langle \mathbf{x}, \mathbf{M}_i\mathbf{x}\rangle + \langle \mathbf{x}, \mathbf{B}_i\mathbf{x}\rangle - 2\langle \psi_i, \mathbf{x}\rangle + z_i^2, \end{aligned} \tag{A.1}$$

where $\mathbf{M}_i = \mathbf{m}_i \otimes \mathbf{m}_i$, $\mathbf{B}_i = \mathrm{E}_\pi[\mathbf{U}_i\mathbf{U}_i^\top] \in \mathrm{PSD}^n$ is the covariance matrix of deviations toward i, and $\psi_i = z_i\mathbf{m}_i$. Summing expression (A.1) over i yields

$$\mathrm{E}_\pi[\|\mathbf{G}\mathbf{x} - \mathbf{z}\|^2] = \langle \mathbf{x}, \mathbf{M}\mathbf{x}\rangle + \langle \mathbf{x}, \mathbf{B}\mathbf{x}\rangle - 2\langle \psi, \mathbf{x}\rangle + \|\mathbf{z}\|^2, \tag{A.2}$$

where $\mathbf{M} = \sum_{i=1}^n \mathbf{M}_i$, $\mathbf{B} = \sum_{i=1}^n \mathbf{B}_i$, and $\psi = \sum_{i=1}^n \psi_i$.

Expression (A.2) decomposes the effect of changing $\mathbf{x}$ into three channels: a mean influence component $\langle \mathbf{x}, \mathbf{M}\mathbf{x}\rangle$, an uncertainty component $\langle \mathbf{x}, \mathbf{B}\mathbf{x}\rangle$, and a target-alignment component $-2\langle \psi, \mathbf{x}\rangle$. In the main text, we combine this with the intervention cost $\|\mathbf{C}^{\frac{1}{2}}(\mathbf{x} - \mathbf{x}^0)\|^2$ to obtain a quadratic representation of the DM's problem.

B Characterization of Property B

We here present an equivalent characterization of Property B in terms of the model primitives and the uncertainty set $\mathcal{B}$, alongside simple sufficient conditions. Recall that $\psi^0 = \mathbf{C}\mathbf{x}^0$ and $\psi = \sum_{i=1}^n z_i\mathbf{m}_i$. For any $\mathbf{x} \in \mathbb{R}^n$, we define the support function of $\mathcal{B}$ evaluated at $\mathbf{x}\mathbf{x}^\top$ as $g(\mathbf{x}) = \max_{\mathbf{B}\in\mathcal{B}} \frac{1}{2}\langle \mathbf{x}, \mathbf{B}\mathbf{x}\rangle$. Note that g is convex and differentiable on the open set $\mathbb{R}_o^n$, which consists of all vectors with strictly non-zero entries:

$$\mathbb{R}_o^n = \{\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n \mid x_i \neq 0 \text{ for all } i = 1, \ldots, n\}.$$

Let R be defined as the image of $\mathbb{R}_o^n$ under the gradient mapping:

$$R = \{(\mathbf{M} + \mathbf{C})\mathbf{x} + \nabla g(\mathbf{x}) \mid \mathbf{x} \in \mathbb{R}_o^n\}.$$

This construction ensures that R is an open set in $\mathbb{R}^n$.

Proposition A1 (Equivalent formulation of Property B) *Assume Property A. Let $\mathbf{x}^*$ denote the unique solution to the robust problem (1). Then, Property B holds if and only if $\psi^0 + \psi \in R$.*

Proof. As shown in Appendix A, omitting constant terms, the robust problem can be restated as:

$$\min_{\mathbf{x}\in\mathbb{R}^n} \left\{ \frac{1}{2}\langle \mathbf{x}, (\mathbf{M}+\mathbf{C})\mathbf{x}\rangle - \langle \psi^0 + \psi, \mathbf{x}\rangle + g(\mathbf{x}) \right\}.$$

Under Property A, $\mathbf{M} \succ 0$. This guarantees that the objective function is strictly convex and that $\mathbf{x}^*$ is its unique minimizer. The first-order optimality condition for this convex program requires:

$$\mathbf{0} \in (\mathbf{M}+\mathbf{C})\mathbf{x}^* - (\psi^0 + \psi) + \partial g(\mathbf{x}^*),$$

which can be rewritten as $\psi^0 + \psi \in (\mathbf{M}+\mathbf{C})\mathbf{x}^* + \partial g(\mathbf{x}^*)$, where $\partial g(\mathbf{x})$ denotes the subdifferential of g at $\mathbf{x}$. Because $\mathbf{x}^*$ is uniquely determined by this optimality condition, and given the definition of R, the proposition immediately follows. ■

Another way to express the equivalent formulation in the proposition is $\psi \in R - \psi^0$, which can be equivalently written in terms of the target vector $\mathbf{z}$:

$$\mathbf{z} \in \mathcal{Z} = \mathrm{E}_\pi[\mathbf{G}^\top]^{-1}(R - \psi^0) = \left\{ \mathrm{E}_\pi[\mathbf{G}^\top]^{-1}\big((\mathbf{M}+\mathbf{C})\mathbf{x} + \nabla g(\mathbf{x}) - \mathbf{C}\mathbf{x}^0\big) \mid \mathbf{x} \in \mathbb{R}^n_o \right\}.$$

Thus, Property B holds whenever the DM's target $\mathbf{z}$ lies in the open set $\mathcal{Z}$.

To illustrate the genericity of this condition, we observe that the set $\mathcal{Z}$ effectively expands to cover the entire space $\mathbb{R}^n$ as the deterministic influence and cost matrix $\mathbf{M}+\mathbf{C}$ increasingly dominates the uncertainty set $\mathcal{B}$.

Proposition A2 (Genericity) *Let $\underline{\lambda} > 0$ denote the minimum eigenvalue of $\mathbf{M}+\mathbf{C}$. Given any bounded set $Q \subset \mathbb{R}^n$, there exists a constant $C > 0$ such that*

$$\mathrm{Leb}(Q \setminus \mathcal{Z}) \leq C/\underline{\lambda}$$

where Leb *denotes the Lebesgue measure. In particular,* $\mathrm{Leb}(Q \setminus \mathcal{Z}) \to 0$ *as* $\underline{\lambda} \to \infty$.

Proof. Let $\Phi : \mathbb{R}^n \to 2^{\mathbb{R}^n}$ be the set-valued mapping defined by $\Phi(\mathbf{x}) = (\mathbf{M}+\mathbf{C})\mathbf{x} + \partial g(\mathbf{x})$. Let $Z_{\text{bound}} = \mathbb{R}^n \setminus \mathbb{R}^n_o = \bigcup_i H_i$ be the union of the n coordinate hyperplanes $H_i = \{\mathbf{x} \in \mathbb{R}^n \mid x_i = 0\}$. The complement of the open set R is exactly the image of this boundary under Φ, meaning $\mathbb{R}^n \setminus R = \Phi(Z_{\text{bound}})$. We are interested in the measure of the target vectors $\mathbf{z} \in Q \setminus \mathcal{Z}$. Recall the linear transformation linking ψ to $\mathbf{z}$: $\mathbf{z} = L(\psi) = \mathrm{E}_\pi[\mathbf{G}^\mathsf{T}]^{-1}\psi$. Let $\varphi = \psi^0 + \psi$, and recall $\mathbf{z} \notin \mathcal{Z} \iff \varphi \notin R$.

Since $\mathbf{z} \in Q$ and Q is bounded, the set of $\psi = L^{-1}(\mathbf{z})$ is also bounded, and thus $\|\varphi\| \leq \|\psi\| + \|\psi^0\| \leq C$ for some $C > 0$. For any target vector that fails Property B, we have $\varphi \notin R$, which means $\varphi \in \Phi(Z_{\text{bound}})$. Thus, there exists $\mathbf{x} \in Z_{\text{bound}}$ such that $\varphi \in \Phi(\mathbf{x})$, meaning $\varphi = (\mathbf{M} + \mathbf{C})\mathbf{x} + \mathbf{B}\mathbf{x}$ for some $\mathbf{B} \in \partial g(\mathbf{x}) \subset \mathcal{B}$. Because $\mathbf{B} \succeq 0$, the minimum eigenvalue of $\mathbf{M} + \mathbf{C} + \mathbf{B}$ is at least $\underline{\lambda}$. Consequently:

$$\underline{\lambda}\|\mathbf{x}\|^2 \leq \langle \mathbf{x}, (\mathbf{M} + \mathbf{C} + \mathbf{B})\mathbf{x} \rangle = \langle \mathbf{x}, \varphi \rangle \leq \|\mathbf{x}\|\|\varphi\| \leq C\|\mathbf{x}\|.$$

This provides an upper bound on the norm of the intervention vector: $\|\mathbf{x}\| \leq C/\underline{\lambda}$.

As $Z_{\text{bound}} = \bigcup_i H_i$, we bound the volume of $\Phi(H_i \cap \{\mathbf{x} \mid \|\mathbf{x}\| \leq C/\underline{\lambda}\})$. The mapping Φ sends $\mathbf{x} \mapsto \mathbf{y} + \mathbf{u}$, where $\mathbf{y} = (\mathbf{M} + \mathbf{C})\mathbf{x}$ lies in a fixed $(n-1)$-dimensional subspace V_i, and $\mathbf{u} = \mathbf{B}\mathbf{x}$. Since $\|\mathbf{u}\| \leq \|\mathbf{B}\|_{\max}\|\mathbf{x}\| \leq C_1(C/\underline{\lambda})$, where $C_1 = \max_{\mathbf{B} \in \mathcal{B}} \|\mathbf{B}\|$ is finite, the set of all such φ is contained within a cylinder around the subspace V_i with an orthogonal radius bounded by $O(1/\underline{\lambda})$. Since $\|\varphi\|$ is bounded by C, this truncated cylinder has a bounded volume that scales proportionally to its radius, yielding an n-dimensional Lebesgue measure of $O(1/\underline{\lambda})$. Summing this volume bound over all n hyperplanes yields $\mathrm{Leb}(\{\varphi \notin R \mid \|\varphi\| \leq C\}) = O(1/\underline{\lambda})$, which in turn yields $\mathrm{Leb}(\{\psi \notin R - \psi^0 \mid \|\psi\| \leq C\}) = O(1/\underline{\lambda})$. Finally, the linear relationship $\mathbf{z} = L(\psi)$ gives $\mathrm{Leb}(Q \setminus \mathcal{Z}) \leq |\det(L)| \cdot O(1/\underline{\lambda})$. ■

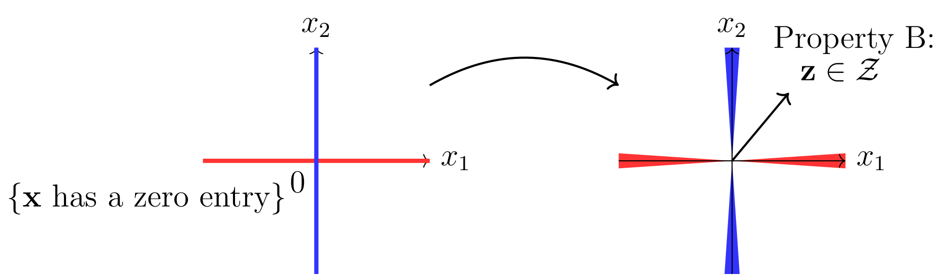

Figure 7: The region $Q \setminus \mathcal{Z}$ is shown as the union of the red and blue areas (in the simple case $\mathbf{M} + \mathbf{C} = \underline{\lambda}\mathrm{Id}$ and $\mathbf{x}^0 = 0$). As $\underline{\lambda} \to \infty$, the region $Q \setminus \mathcal{Z}$ shrinks.

C Proofs

Proof of Theorem 1. We prove the three parts sequentially as follows.

Part (a). Recall $f(\mathbf{x}, \mathbf{B})$ from expression (3) in the main text:

$$f(\mathbf{x}, \mathbf{B}) \;=\; \frac{1}{2}\Big(\langle \mathbf{x}, (\mathbf{M} + \mathbf{C})\mathbf{x}\rangle + \langle \mathbf{x}, \mathbf{B}\mathbf{x}\rangle - 2\langle \psi^0 + \psi, \mathbf{x}\rangle\Big).$$

Under Property A, $\mathbf{M}$ is positive definite, so for every $\mathbf{B} \in \mathcal{B}$ the function $\mathbf{x} \mapsto f(\mathbf{x}, \mathbf{B})$ is strictly convex and coercive. For every $\mathbf{x} \in \mathbb{R}^n$, the function $\mathbf{B} \mapsto f(\mathbf{x}, \mathbf{B})$ is linear. The set $\mathcal{B}$ is convex and compact, and $\mathbb{R}^n$ is convex. Thus, by Sion's minimax theorem, the minimax equality (4) in the theorem holds, and so there exists a saddle point $(\mathbf{x}^*, \mathbf{B}^*)$ such that $f(\mathbf{x}^*, \mathbf{B}) \le f(\mathbf{x}^*, \mathbf{B}^*) \le f(\mathbf{x}, \mathbf{B}^*)$ for all $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{B} \in \mathcal{B}$.

At a saddle point, $\mathbf{x}^*$ minimizes $\mathbf{x} \mapsto f(\mathbf{x}, \mathbf{B}^*)$. Differentiating in $\mathbf{x}$ yields $(\mathbf{M} + \mathbf{B}^* + \mathbf{C})\mathbf{x}^* = \psi^0 + \psi$, which is (5) in the theorem. Uniqueness of $\mathbf{x}^*$ follows because $\mathbf{x} \mapsto \max_{\mathbf{B} \in \mathcal{B}} f(\mathbf{x}, \mathbf{B})$ is strictly convex: it is a pointwise maximum of strictly convex functions in $\mathbf{x}$, and strict convexity is inherited because the quadratic term $\frac{1}{2}\langle \mathbf{x}, \mathbf{M}\mathbf{x}\rangle$ is strictly convex. Therefore, Part (a) is proven.

Part (b). Fix $\mathbf{x} \in \mathbb{R}^n$ and $i \in N$. Nature's row-wise problem is to maximize $\langle \mathbf{x}, \mathbf{B}_i\mathbf{x}\rangle$ over $\mathbf{B}_i \in \mathcal{B}_i$, and these problems are independent across i. Let $\mathbf{D}_i = \operatorname{diag}(\mathbf{v}_{i1}, \ldots, \mathbf{v}_{in})$ and write $\mathbf{y} = \mathbf{D}_i\mathbf{x}$, so $y_j = \mathbf{v}_{ij}x_j$. Any $\mathbf{B}_i \in \mathcal{B}_i$ can be written as $\mathbf{B}_i = \mathbf{D}_i\mathbf{R}_i\mathbf{D}_i$ for some correlation matrix $\mathbf{R}_i \succeq 0$ with $(\mathbf{R}_i)_{jj} = 1$ for all j. Then $\langle \mathbf{x}, \mathbf{B}_i\mathbf{x}\rangle = \langle \mathbf{y}, \mathbf{R}_i\mathbf{y}\rangle$.

Using $|(\mathbf{R}_i)_{jk}| \le 1$ and symmetry, we have

$$\langle \mathbf{y}, \mathbf{R}_i\mathbf{y}\rangle = \sum_{j=1}^{n} y_j^2 + 2\sum_{1 \le j < k \le n} (\mathbf{R}_i)_{jk} y_j y_k \le \sum_{j=1}^{n} y_j^2 + 2\sum_{1 \le j < k \le n} |y_j y_k| = \Big(\sum_{j=1}^{n} |y_j|\Big)^2.$$

If $\mathbf{x}$ has no zero entries, then $y_j \neq 0$ for all j and equality holds if and only if $(\mathbf{R}_i)_{jk} = s(y_j)s(y_k)$ for all j, k, which forces $\mathbf{R}_i = \mathbf{s}(\mathbf{y}) \otimes \mathbf{s}(\mathbf{y})$, where $\mathbf{s}(\mathbf{y}) = (s(y_1), \ldots, s(y_n))^\mathsf{T}$. Therefore, the unique maximizer is

$$\mathbf{B}_{i,\mathrm{BR}}(\mathbf{x}) = \mathbf{D}_i\big(\mathbf{s}(\mathbf{y}) \otimes \mathbf{s}(\mathbf{y})\big)\mathbf{D}_i \;=\; \mathbf{q}_i(\mathbf{x}) \otimes \mathbf{q}_i(\mathbf{x}),$$

where $\mathbf{q}_i(\mathbf{x})$ is defined in (6) and we used $s(y_j) = s(\mathbf{v}_{ij}x_j) = s(x_j)$ because $\mathbf{v}_{ij} > 0$. This proves the rank-1 form and uniqueness of the row-wise best response when $\mathbf{x}$ has no zero entries. Aggregating over i yields (7) in the theorem.

If $\mathbf{x}$ has a zero entry, then some $y_j = 0$ and the sign restriction on $(\mathbf{R}_i)_{jk}$ is not pinned down for pairs involving j, so the maximizer need not be unique; any $\mathbf{R}_i$ satisfying $(\mathbf{R}_i)_{jk} = s(y_j)s(y_k)$ whenever $y_j y_k \neq 0$ attains the same value.

For part (b), Property B guarantees that $\mathbf{x}^*$ has no zero entries, so the row-wise best-response formula applies at $\mathbf{x} = \mathbf{x}^*$. Combined with the saddle-point property, $\mathbf{B}^* = \mathbf{B}_{\mathrm{BR}}(\mathbf{x}^*) = \sum_{i=1}^{n} \mathbf{q}_i(\mathbf{x}^*) \otimes \mathbf{q}_i(\mathbf{x}^*)$. Substituting into (5) yields (8).

Part (c). Part (c) follows from the uniqueness statement in the above rank-1 property: under Property B, $\mathbf{x}^*$ has no zero entries, so $\mathbf{B}_{\text{BR}}(\mathbf{x}^*)$ is unique. Therefore, the worst-case scenario $\mathbf{B}^*$ is unique. ■

Proof of Proposition 1. Consider a d-regular network G_d with n agents, uniform mean link weight $m > 0$, and uniform variance $v^2 \geq 0$. We restrict attention to a symmetric intervention $\mathbf{x} = x\mathbf{1}_n$, a uniform target $z_i = 1$ for all i, and a scalar cost $c > 0$. We also assume $\mathbf{x}^0 = \mathbf{0}$ and $\mathbf{C} = c\mathbf{I}$. We derive each term step by step as follows.

Step 1: Mean influence term. For each agent i, the mean influence vector $\mathbf{m}_i \in \mathbb{R}^n$ has entries $m_{ij} = m$ for all j in i's neighborhood N_i (with $|N_i| = d$) and zero otherwise. Thus, $\mathbf{m}_i^\mathsf{T}\mathbf{x} = dmx$ for all i. Since $\mathbf{M} = \sum_{i=1}^n \mathbf{m}_i \otimes \mathbf{m}_i$,

$$\mathbf{x}^\mathsf{T}\mathbf{M}\mathbf{x} = \sum_{i=1}^n (\mathbf{m}_i^\mathsf{T}\mathbf{x})^2 = n(dmx)^2 = nd^2m^2x^2.$$

Step 2: Worst-case uncertainty term. By Theorem 1-(b), Nature's worst-case row covariance is $\mathbf{B}_i^* = \mathbf{q}_i(\mathbf{x}) \otimes \mathbf{q}_i(\mathbf{x})$, where $\mathbf{q}_i(\mathbf{x}) = (s(x_1)\mathbf{v}_{i1}, \ldots, s(x_n)\mathbf{v}_{in})^\mathsf{T}$. Under a symmetric positive intervention $\mathbf{x} = x\mathbf{1}_n$ with $x > 0$, all signs are positive. Since $\mathbf{v}_{ij} = v$ for $j \in N_i$ and $\mathbf{v}_{ij} = 0$ otherwise, we have $\mathbf{q}_i(\mathbf{x})^\mathsf{T}\mathbf{x} = \sum_{j \in N_i} v \cdot x = dvx$. Consequently,

$$\mathbf{x}^\mathsf{T}\mathbf{B}^*\mathbf{x} = \sum_{i=1}^n (\mathbf{q}_i(\mathbf{x})^\mathsf{T}\mathbf{x})^2 = n(dvx)^2 = nd^2v^2x^2.$$

Step 3: Intervention cost term. The cost of deviating from $\mathbf{x}^0 = \mathbf{0}$ is $\frac{1}{2}\|\mathbf{C}^{\frac{1}{2}}(\mathbf{x} - \mathbf{x}^0)\|^2 = \frac{c}{2}\|\mathbf{x}\|^2 = \frac{cn}{2}x^2$.

Step 4: Target-alignment term. The leverage vector is $\psi = \sum_{i=1}^n z_i\mathbf{m}_i = \sum_{i=1}^n \mathbf{m}_i$. Since each agent j appears in exactly d neighborhoods, $\psi_j = dm$ for all j. Hence, $\langle \psi, \mathbf{x} \rangle = dm \cdot nx = ndmx$.

Combining Steps 1-4 and defining $\Omega = m^2 + v^2$, the DM's robust objective becomes

$$\frac{1}{2}\Big(nd^2m^2x^2 + nd^2v^2x^2 + cnx^2\Big) - ndmx = \frac{n}{2}\big(d^2\Omega + c\big)x^2 - ndmx,$$

which is expression (9). The first-order condition yields $n(d^2\Omega + c)x = ndm$, which implies that $x^*(d) = \frac{dm}{d^2(m^2+v^2)+c}$.

We now present the inverted-U shape of the optimal intervention $x^*(d)$. Treating $x^*(d)$ as a function of the continuous variable $d > 0$, it follows that

$$\frac{\partial x^*}{\partial d} = \frac{m(c - d^2\Omega)}{(d^2\Omega + c)^2}.$$

Setting the numerator to zero gives $d^* = \sqrt{\frac{c}{\Omega}} = \sqrt{\frac{c}{(m^2+v^2)}}$. Since $\frac{\partial^2 x^*}{\partial d^2}\Big|_{d=d^*} < 0$, the intervention intensity $x^*(d)$ is maximized at d^*, confirming the inverted-U shape. ■

Proof of Proposition 2. Consider the core-periphery network $CP_{n,k}$ with k core agents and $n-k$ periphery agents, uniform mean influence $m > 0$, uniform variance $v^2 \geq 0$, target $\mathbf{z} = \mathbf{1}$, cost $\mathbf{C} = c\mathbf{I}$, and reference $\mathbf{x}^0 = \mathbf{0}$. We use the block-symmetric intervention $\mathbf{x} = (x_c \mathbf{1}_k^\mathsf{T}, x_p \mathbf{1}_{n-k}^\mathsf{T})^\mathsf{T}$. We derive each term step by step as follows.

Step 1: Mean influence term. The mean adjacency matrix $\overline{\mathbf{G}}$ has $\overline{\mathbf{G}}_{ij} = m$ whenever agents i and j are linked and zero otherwise. In addition, since $\overline{\mathbf{G}}$ is symmetric, $\mathbf{M} = \overline{\mathbf{G}}^\mathsf{T}\overline{\mathbf{G}} = \overline{\mathbf{G}}^2$. For each core agent $i \leq k$, i is linked to $k-1$ other core agents and $n-k$ periphery agents. Hence, it follows that $y_c = (\overline{\mathbf{G}}\mathbf{x})_i = m\big[(k-1)x_c + (n-k)x_p\big]$. On the contrary, for each periphery agent $i > k$, agent i is linked only to the k core agents, which implies that $y_p = (\overline{\mathbf{G}}\mathbf{x})_i = m(kx_c)$. Consequently, we have

$$\begin{aligned}\mathbf{x}^\mathsf{T}\mathbf{M}\mathbf{x} &= \|\overline{\mathbf{G}}\mathbf{x}\|^2 = k \cdot y_c^2 + (n-k) \cdot y_p^2 \\ &= m^2\big[k\big((k-1)x_c + (n-k)x_p\big)^2 + (n-k)(kx_c)^2\big].\end{aligned}$$

Step 2: Worst-case uncertainty term. By Theorem 1-(b), we obtain $\mathbf{x}^\mathsf{T}\mathbf{B}^*\mathbf{x} = \sum_{i=1}^n (\sum_{j\in N_i} \mathbf{v}_{ij}|x_j|)^2$. Under uniform v and a positive intervention (guess and verify), this takes the same block form as Step 1 with v replacing m:

$$\mathbf{x}^\mathsf{T}\mathbf{B}^*\mathbf{x} = v^2\big[k\big((k-1)x_c + (n-k)x_p\big)^2 + (n-k)(kx_c)^2\big].$$

Step 3: Combined quadratic and cost. Defining $\Omega = m^2 + v^2$ and combining Step 1 and Step 2, we have

$$\mathbf{x}^\mathsf{T}(\mathbf{M} + \mathbf{B}^*)\mathbf{x} = \Omega\big[k\big((k-1)x_c + (n-k)x_p\big)^2 + (n-k)(kx_c)^2\big].$$

Similarly, the intervention cost is $\frac{c}{2}(kx_c^2 + (n-k)x_p^2)$.

Step 4: Target-alignment term. For core agent $j \leq k$: agent j appears in the neighborhood of all other $k-1$ core agents and all $n-k$ periphery agents, so $\psi_j = (k-1+n-k)m = m(n-1) =: \psi_c$. For periphery agent $j > k$: agent j appears only in the neighborhoods of the k core agents, so $\psi_j = km =: \psi_p$. Hence,

$$\langle \psi, \mathbf{x} \rangle = k\psi_c x_c + (n-k)\psi_p x_p = km(n-1)x_c + (n-k)mkx_p.$$

We now characterize the optimal intervention. Differentiating the full objective with respect to x_c and dividing by k, using $(k-1)^2 + k(n-k) = (n-2)k+1$, it follows that

$$\big[\Omega\big((n-2)k+1\big) + c\big]x_c + \Omega(k-1)(n-k)x_p = m(n-1).$$

Differentiating with respect to x_p and dividing by $(n-k)$ yields

$$\Omega k(k-1)x_c + \big[\Omega k(n-k) + c\big]x_p = mk.$$

By combining the above two equations, we have the following characterization in matrix form:

$$\underbrace{\begin{bmatrix} \Omega((n-2)k+1) + c & \Omega(k-1)(n-k) \\ \Omega k(k-1) & \Omega k(n-k) + c \end{bmatrix}}_{=:\mathbf{A}} \begin{bmatrix} x_c \\ x_p \end{bmatrix} = m \begin{bmatrix} n-1 \\ k \end{bmatrix}.$$

The coefficient matrix $\mathbf{A}$ has positive diagonal entries. Its determinant is

$$\begin{aligned} \det(\mathbf{A}) &= \big[\Omega((n-2)k+1) + c\big]\big[\Omega k(n-k) + c\big] - \Omega^2 k(k-1)^2(n-k) \\ &= c^2 + c\Omega\big[(n-2)k + 1 + k(n-k)\big] + \Omega^2 k(n-k)\big[(n-2)k+1-(k-1)^2\big]. \end{aligned}$$

We compute $(n-2)k+1-(k-1)^2 = k(n-k)$, so the last term equals $\Omega^2 k^2(n-k)^2 > 0$. Hence, its determinant is strictly positive. Consequently, by Cramer's rule, we obtain

$$x_c^* = \frac{m\big[\Omega k(n-k)^2 + c(n-1)\big]}{\det(\mathbf{A})} > 0, \quad x_p^* = \frac{mk\big[\Omega(n-k) + c\big]}{\det(\mathbf{A})} > 0,$$

where the numerator for x_c^* follows from $(n-1)-(k-1) = n-k$, and that for x_p^* from $(n-2)k+1-(n-1)(k-1) = n-k$. This verifies the initial sign guess. ■

Proof of Proposition 3. Consider the directed bipartite network with k upstream agents $U = \{1, \dots, k\}$ and $n-k$ downstream agents $D = \{k+1, \dots, n\}$. Each upstream agent $j \in U$ influences each downstream agent $i \in D$ with mean $m > 0$ and variance $v^2 \geq 0$. Each downstream agent i has a self-loop with mean $\mu > 0$ and variance $v_D^2 \geq 0$. Influence flows only from U to D. The target is $z_i = 1$ for $i \in D$ and $z_i = 0$ for $i \in U$, with cost $\mathbf{C} = c\mathbf{I}$ and reference $\mathbf{x}^0 = \mathbf{0}$. We use the block intervention $\mathbf{x} = (x_U \mathbf{1}_k^\mathsf{T}, x_D \mathbf{1}_{n-k}^\mathsf{T})^\mathsf{T}$, so Property A holds for the reduced mean-influence matrix.

Step 1: Mean influence structure. For upstream agents $i \in U$, since no influence flows to upstream, $\mathbf{m}_i = \mathbf{0}$, hence $\mathbf{M}_i = \mathbf{0}$. For downstream agent $i \in D$: $\mathbf{m}_i \in \mathbb{R}^n$ has entries $m_{ij} = m$ for $j \in U$, $m_{ii} = \mu$ (self-loop), and $m_{ij} = 0$ for $j \in D \setminus \{i\}$. Thus,

$$\mathbf{m}_i^\mathsf{T}\mathbf{x} = kmx_U + \mu x_D \quad \text{for all } i \in D,$$

and $\mathbf{x}^\mathsf{T}\mathbf{M}\mathbf{x} = \sum_{i\in D}(\mathbf{m}_i^\mathsf{T}\mathbf{x})^2 = (n-k)(kmx_U + \mu x_D)^2$.

Step 2: Worst-case uncertainty. For upstream agents $i \in U$, $\mathbf{v}_{ij} = 0$ for all j, so $\mathbf{B}_i^* = \mathbf{0}$. For downstream agent $i \in D$, the variance profile has $\mathbf{v}_{ij} = v$ for $j \in U$, $\mathbf{v}_{ii} = v_D$ (self-loop), and $\mathbf{v}_{ij} = 0$ for $j \in D \setminus \{i\}$. Under a positive intervention (guess and verify), by Theorem 1-(b), we have $\mathbf{q}_i(\mathbf{x})^\mathsf{T}\mathbf{x} = kv \cdot x_U + v_D \cdot x_D$ for all $i \in D$. Hence, it follows that

$$\mathbf{x}^\mathsf{T}\mathbf{B}^*\mathbf{x} = \sum_{i\in D}(kvx_U + v_Dx_D)^2 = (n-k)(kvx_U + v_Dx_D)^2.$$

Step 3: Expansion. Combining Steps 1 and 2 and expanding, we obtain

$$\begin{aligned}\mathbf{x}^\mathsf{T}(\mathbf{M}+\mathbf{B}^*)\mathbf{x} &= (n-k)\big[(kmx_U + \mu x_D)^2 + (kvx_U + v_Dx_D)^2\big] \\ &= (n-k)\big[k^2\Omega\, x_U^2 + 2k\Omega_{UD}\, x_Ux_D + \Omega_D\, x_D^2\big],\end{aligned}$$

where $\Omega = m^2 + v^2$, $\Omega_{UD} = m\mu + vv_D$, and $\Omega_D = \mu^2 + v_D^2$.

Step 4: Target-alignment and cost. Since $z_i = 0$ for $i \in U$, the leverage vector satisfies $\psi_j = \sum_{i\in D} z_i m_{ij}$. For $j \in U$: $\psi_j = (n-k)m$. For $j \in D$: $\psi_j = \mu$ (from the self-loop of the unique $i = j$). Hence,

$$\langle \psi, \mathbf{x}\rangle = k(n-k)m\, x_U + (n-k)\mu\, x_D = (n-k)\big[km\, x_U + \mu\, x_D\big].$$

The cost term is $\frac{c}{2}[kx_U^2 + (n-k)x_D^2]$.

We observe that the full objective is

$$\begin{aligned}&\frac{n-k}{2}\big[k^2\Omega\, x_U^2 + 2k\Omega_{UD}\, x_Ux_D \\ &\quad + \Omega_D\, x_D^2\big] + \frac{c}{2}\big[kx_U^2 + (n-k)x_D^2\big] - (n-k)\big[km\, x_U + \mu\, x_D\big],\end{aligned}$$

which is expression (10) in the main text. Differentiating with respect to x_U and x_D and rearranging yields the following linear system:

$$\underbrace{\begin{bmatrix} k^2(n-k)\Omega + kc & k(n-k)\Omega_{UD} \\ k(n-k)\Omega_{UD} & (n-k)\Omega_D + (n-k)c \end{bmatrix}}_{=:\mathbf{A}} \begin{bmatrix} x_U^* \\ x_D^* \end{bmatrix} = (n-k)\begin{bmatrix} km \\ \mu \end{bmatrix}.$$

In order to verify positive definiteness of $\mathbf{A}$, we apply Sylvester's criterion. First, the $(1, 1)$ entry $k^2(n-k)\Omega + kc = k[k(n-k)\Omega + c] > 0$. Second, for the determinant, expanding and collecting terms yields

$$\begin{aligned}\det(\mathbf{A}) &= \big[k^2(n-k)\Omega + kc\big]\big[(n-k)\Omega_D + (n-k)c\big] - k^2(n-k)^2\Omega_{UD}^2 \\ &= k(n-k)\Big\{k(n-k)\big[\Omega\Omega_D - \Omega_{UD}^2\big] + c\big[k(n-k)\Omega + \Omega_D + c\big]\Big\}.\end{aligned}$$

Now $\Omega\Omega_D - \Omega_{UD}^2 = (m^2 + v^2)(\mu^2 + v_D^2) - (m\mu + vv_D)^2 = (mv_D - v\mu)^2 \geq 0$, and all remaining terms are strictly positive. Therefore, we have $\det(\mathbf{A}) > 0$, so $\mathbf{A}$ is positive definite. By Cramer's rule,

$$x_U^* = \frac{k(n-k)^2[mc + v_D(mv_D - \mu v)]}{\det(\mathbf{A})}, \quad x_D^* = \frac{k(n-k)[k(n-k)v(\mu v - mv_D) + \mu c]}{\det(\mathbf{A})}.$$

For c sufficiently large (specifically, $c > |mv_D - \mu v| \cdot \max\{v_D/m,\ k(n-k)v/\mu\}$), both x_U^* and x_D^* are strictly positive, verifying the sign guess.

Comparative statics. Write $\det \mathbf{A} = k(n-k)D$ where $D = k(n-k)\Delta^2 + c[k(n-k)\Omega + \Omega_D + c]$ with $\Delta := mv_D - \mu v$. Then,

$$x_U^* = \frac{(n-k)[mc + v_D\Delta]}{D}, \qquad x_D^* = \frac{\mu c - k(n-k)v\Delta}{D}.$$

When $v_D = 0$, we have $\Delta = -\mu v$, so $x_U^* = (n-k)mc/D$ and $x_D^* = \mu[c + k(n-k)v^2]/D$. The ratio simplifies to $\frac{x_U^*}{x_D^*} = \frac{(n-k)m/\mu}{1+k(n-k)v^2/c}$, which is strictly decreasing in v. When $v = 0$, we have $\Delta = mv_D$, so $x_U^* = (n-k)m(c + v_D^2)/D$ and $x_D^* = \mu c/D$. The ratio is

$$\frac{x_U^*}{x_D^*} = \frac{(n-k)m}{\mu}\Big(1 + \frac{v_D^2}{c}\Big),$$

which is strictly increasing in v_D.

Now suppose that both v_D and v are strictly positive. Differentiating x_U^* with respect to v at $v = 0$ gives

$$\frac{\partial x_U^*}{\partial v}\bigg|_{v=0} = \frac{(n-k)\mu v_D\big[-D(0) + 2k(n-k)m^2(c + v_D^2)\big]}{D(0)^2}.$$

The bracket is negative iff $D(0) > 2k(n-k)m^2(c + v_D^2)$. Direct calculation yields

$$D(0) - 2k(n-k)m^2(c + v_D^2) = [c - k(n-k)m^2](v_D^2 + c) + c\mu^2,$$

which is strictly positive whenever $c \geq k(n-k)m^2$. Similarly, differentiating x_D^* at $v = 0$ gives

$$\frac{\partial x_D^*}{\partial v}\bigg|_{v=0} = \frac{k(n-k)mv_D\big[-D(0) + 2\mu^2 c\big]}{D(0)^2}.$$

We have $D(0) - 2\mu^2 c = k(n-k)m^2(v_D^2 + c) + c(c + v_D^2 - \mu^2)$, which is strictly positive whenever $c \geq \mu^2$.

Therefore, when $v_D > 0$ and $c \geq \max\{k(n-k)m^2,\ \mu^2\}$, both $\partial x_U^*/\partial v|_{v=0} < 0$ and $\partial x_D^*/\partial v|_{v=0} < 0$. ■

Proof of Proposition 4. We state and prove a generalized version.[16] Recall that PD^n and PSD^n denote the sets of all $n \times n$ real symmetric positive definite and positive semi-definite matrices.

Proposition A3 *Let* $\mathbf{B} \in \mathrm{PD}^{\mathrm{n}}$*, I be a proper (possibly empty) subset of $N = \{1, ..., n\}$, $b_i \in \mathbb{R}$ for all $i \in I$, and $b_{n+1} > 0$. Then, for any intervention $\overline{\mathbf{x}} \in \mathbb{R}^{n+1}$ with $\overline{x}_j \overline{x}_{n+1} \neq 0$ for some $j \in \{1, ..., n\} \setminus I$, the problem*

$$\text{maximize } \langle \overline{\mathbf{x}}, \overline{\mathbf{B}}\overline{\mathbf{x}} \rangle \quad \text{over} \quad \{\overline{\mathbf{B}} \in \overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B},b} \,|\, \overline{\mathbf{B}}_{i,n+1} = b_i \text{ for all } i \in I\} \tag{C.1}$$

has a unique solution (unless the latter feasible set is empty).

For a given matrix $\mathbf{B} \in \mathbb{R}^{n\times n}$, let $\overline{\mathbf{B}} \in \mathbb{R}^{(n+1)\times(n+1)}$ denote a symmetric and extended matrix of $\mathbf{B}$ having $\mathbf{B}$ as its principal submatrix. For each $\mathbf{v} = (\mathbf{v}_1, ..., \mathbf{v}_n)^\mathsf{T} \in \mathbb{R}^n$, let $\overline{\mathbf{v}} = (\overline{\mathbf{v}}_1, \ldots, \overline{\mathbf{v}}_n, \overline{\mathbf{v}}_{n+1})^\mathsf{T} \in \mathbb{R}^{n+1}$ denote the extended vector of $\mathbf{v}$ having $\mathbf{v}$ as its subvector; that is, $\overline{\mathbf{v}}_k = \mathbf{v}_k$ for all $1 \leq k \leq n$. By abusing notation, we denote $\overline{\mathbf{v}}$ by $\overline{\mathbf{v}} = (\mathbf{v}^\mathsf{T}, \mathbf{v}_{n+1})^\mathsf{T}$. For a given matrix $\mathbf{A}$, we let $\mathrm{Null}(\mathbf{A})$ denote the Null space of $\mathbf{A}$; that is, $\mathrm{Null}(\mathbf{A}) = \{\mathbf{x} \,|\, \mathbf{A}\mathbf{x} = \mathbf{0}\}$.

Let V be an affine subspace of $\mathbb{R}^d$ that inherits the usual topology of $\mathbb{R}^d$; that is, $A \subset V$ is open in V if and only if there exists an open set $\overline{A}$ in $\mathbb{R}^d$ such that $A = \overline{A} \cap V$. We say that a compact convex set $A \subset V$ is strictly convex in V if A has nonempty interior in V, and for any distinct $\mathbf{x}$ and $\mathbf{x}'$ in A, $\frac{\mathbf{x}+\mathbf{x}'}{2}$ is also contained in the interior of A.

Given $\mathbf{B} \in \mathrm{PSD}^n$ and $b > 0$, we define $\overline{\mathcal{B}}_{\mathbf{B},b}$ as the set of all extended matrices of $\mathbf{B}$ having the $(n+1)$-th row and $(n+1)$-th column entry by b:

$$\overline{\mathcal{B}}_{\mathbf{B},b} = \{\overline{\mathbf{B}} \in \mathbb{R}^{(n+1)\times(n+1)} \,|\, \overline{\mathbf{B}} \text{ is symmetric, } \overline{\mathbf{B}}_{ij} = \mathbf{B}_{ij} \text{ for all } 1 \leq i,j \leq n, \overline{\mathbf{B}}_{(n+1)(n+1)} = b\}.$$

We define $\overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B},b} = \overline{\mathcal{B}}_{\mathbf{B},b} \cap \mathrm{PSD}^{n+1}$ and $\overline{\mathcal{B}}^{\mathrm{PD}}_{\mathbf{B},b} = \overline{\mathcal{B}}_{\mathbf{B},b} \cap \mathrm{PD}^{n+1}$.

Proposition A3 will be shown as a consequence of the following lemma.

Lemma A1 *For any $\mathbf{B} \in \mathrm{PD}^n$ and $b > 0$, the set*

$$\mathcal{X}^{\mathrm{PSD}}_{\mathbf{B},b} = \{\mathbf{b} \in \mathbb{R}^n \,|\, \mathbf{b}_i = \overline{\mathbf{B}}_{i,n+1} \text{ for all } i = 1, ..., n \text{ for some } \overline{\mathbf{B}} \in \overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B},b}\} \tag{C.2}$$

[16] Proposition 4 is clearly a special case of Proposition A3 with $I = \emptyset$. We shall prove Proposition A3.

is a strictly convex and compact subset of $\mathbb{R}^n$*.*

Proof. It is clear that $\mathcal{X}^{\mathrm{PSD}}_{\mathbf{B},b}$ is compact and convex. For $\mathbf{B} \in \mathrm{PD}^n$ and $b > 0$, by Sylvester's criterion (Meyer, 2010), $\overline{\mathbf{B}} \in \overline{\mathcal{B}}_{\mathbf{B},b}$ is positive semi-definite if and only if $\det(\overline{\mathbf{B}}) \geq 0$. Moreover, $\overline{\mathbf{B}} \in \overline{\mathcal{B}}_{\mathbf{B},b}$ is positive definite if and only if $\det(\overline{\mathbf{B}}) > 0$. Thus,

$$\mathcal{X}^{\mathrm{PSD}}_{\mathbf{B},b} = \{\mathbf{b} \in \mathbb{R}^n \,|\, \mathbf{b}_i = \overline{\mathbf{B}}_{i,n+1} \text{ for all } i \in N \text{ for some } \overline{\mathbf{B}} \in \overline{\mathcal{B}}_{\mathbf{B},b} \text{ with } \det(\overline{\mathbf{B}}) \geq 0\},$$

and the interior of the set $\mathcal{X}^{\mathrm{PSD}}_{\mathbf{B},b}$ is given by

$$\begin{aligned}\mathcal{X}^{\mathrm{PD}}_{\mathbf{B},b} &= \{\mathbf{b} \in \mathbb{R}^n \,|\, \mathbf{b}_i = \overline{\mathbf{B}}_{i,n+1} \text{ for all } i \in N \text{ for some } \overline{\mathbf{B}} \in \overline{\mathcal{B}}_{\mathbf{B},b} \text{ with } \det(\overline{\mathbf{B}}) > 0\} \\ &= \{\mathbf{b} \in \mathbb{R}^n \,|\, \mathbf{b}_i = \overline{\mathbf{B}}_{i,n+1} \text{ for all } i = 1, ..., n \text{ for some } \overline{\mathbf{B}} \in \overline{\mathcal{B}}^{\mathrm{PD}}_{\mathbf{B},b}\},\end{aligned}$$

which is nonempty because $\mathbf{0} \in \mathcal{X}^{\mathrm{PD}}_{\mathbf{B},b}$. Let $\mathbf{b}^1, \mathbf{b}^2 \in \mathcal{X}^{\mathrm{PSD}}_{\mathbf{B},b}$ with $\mathbf{b}^1 \neq \mathbf{b}^2$. We need to show $\frac{\mathbf{b}^1+\mathbf{b}^2}{2} \in \mathcal{X}^{\mathrm{PD}}_{\mathbf{B},b}$. For $j = 1, 2$, let $\overline{\mathbf{B}}^j \in \overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B},b}$ such that $\mathbf{b}^j_i = \overline{\mathbf{B}}^j_{i,n+1}$ for all $i = 1, ..., n$. Then, the claim $\frac{\mathbf{b}^1+\mathbf{b}^2}{2} \in \mathcal{X}^{\mathrm{PD}}_{\mathbf{B},b}$ is equivalent to $\det(\frac{\overline{\mathbf{B}}^1+\overline{\mathbf{B}}^2}{2}) > 0$. By a way of contradiction, suppose that this is false. Then, $\det(\frac{\overline{\mathbf{B}}^1+\overline{\mathbf{B}}^2}{2}) = 0$, i.e., $\frac{\overline{\mathbf{B}}^1+\overline{\mathbf{B}}^2}{2}$ is not positive definite, which is the case if and only if $\mathrm{Null}(\overline{\mathbf{B}}^1) \cap \mathrm{Null}(\overline{\mathbf{B}}^2) \neq \{\mathbf{0}\}$. For any $\overline{\mathbf{z}} \in \mathrm{Null}(\overline{\mathbf{B}}^1) \cap \mathrm{Null}(\overline{\mathbf{B}}^2)$ with $\overline{\mathbf{z}} \neq \mathbf{0}$, we claim that $\overline{\mathbf{z}}_{n+1} = 0$. To prove this, suppose $\overline{\mathbf{z}}_{n+1} \neq 0$ and $\overline{\mathbf{B}}^j\overline{\mathbf{z}} = \mathbf{0}$ for $j = 1, 2$. Since $\overline{\mathbf{B}}^j\overline{\mathbf{z}} = \sum_{k=1}^{n+1} \overline{\mathbf{z}}_k \overline{\mathbf{B}}^j_k$, where $\overline{\mathbf{B}}^j_k$ denotes the kth column of $\overline{\mathbf{B}}^j$, the equation $\sum_{k=1}^{n+1} \overline{\mathbf{z}}_k \overline{\mathbf{B}}^j_k = \mathbf{0} \in \mathbb{R}^{n+1}$ in particular yields, by ignoring the last $(n+1)$th row, $\sum_{k=1}^{n} \overline{\mathbf{z}}_k \mathbf{B}_k + \overline{\mathbf{z}}_{n+1}\mathbf{b}^j = \mathbf{0} \in \mathbb{R}^n$, $j = 1, 2$. However, this yields $\mathbf{b}^1 = -\sum_{k=1}^{n} \frac{\overline{\mathbf{z}}_k}{\overline{\mathbf{z}}_{n+1}} \mathbf{B}_k = \mathbf{b}^2$, a contradiction. Hence $\overline{\mathbf{z}}_{n+1} = \mathbf{0}$, but then $\sum_{k=1}^{n} \overline{\mathbf{z}}_k \mathbf{B}_k = \mathbf{0}$ implies $\overline{\mathbf{z}} = \mathbf{0}$ due to the assumption $\mathbf{B} \in \mathrm{PD}^n$. We conclude $\mathrm{Null}(\overline{\mathbf{B}}^1) \cap \mathrm{Null}(\overline{\mathbf{B}}^2) = \{\mathbf{0}\}$, which yields $\det(\frac{\overline{\mathbf{B}}^1+\overline{\mathbf{B}}^2}{2}) > 0$, as desired. ■

We now prove the theorem. As the only unknown entries in $\overline{\mathbf{B}}$ are $(\overline{\mathbf{B}}_{i,n+1})_{i \in \{1,...,n\} \setminus I}$, the problem (C.1) is equivalent to $\max_{\overline{\mathbf{B}} \in \overline{\mathcal{B}}^{\mathrm{PSD}}_{\mathbf{B},b}} \sum_{i \in \{1,...,n\} \setminus I} \overline{\mathbf{x}}_i \overline{\mathbf{x}}_{n+1} \overline{\mathbf{B}}_{i,n+1}$. The set of feasible variables $(\overline{\mathbf{B}}_{i,n+1}) \in \mathbb{R}^{n-|I|}$ is equal to the projection of the slice set $\mathcal{X}^{\mathrm{PSD}}_{\mathbf{B},b} \cap \{\mathbf{x} \in \mathbb{R}^n \,|\, \mathbf{x}_i = b_i \text{ for all } i \in I\}$ onto $\{(\mathbf{x}_i)_{i \in \{1,...,n\} \setminus I}\} \cong \mathbb{R}^{n-|I|}$. Since the objective function is nonzero and linear in the variable $(\overline{\mathbf{B}}_{i,n+1})$, and any slice of a strictly convex set is also strictly convex, the theorem follows from Lemma A1. ■

Proof of Proposition 5. We make repeated use of the two first-order conditions. The naive intervention $\mathbf{x}^N$ minimizes $f(\cdot, \mathbf{0})$, hence

$$(\mathbf{M} + \mathbf{C})\mathbf{x}^N \;=\; \psi^0 + \psi.$$

The robust intervention $\mathbf{x}^*$ satisfies the saddle-point first-order condition (5),

$$(\mathbf{M}+\mathbf{B}^*+\mathbf{C})\mathbf{x}^* \;=\; \psi^0+\psi.$$

Subtracting these yields the identity

$$(\mathbf{M}+\mathbf{C})(\mathbf{x}^N-\mathbf{x}^*) \;=\; \mathbf{B}^*\mathbf{x}^*, \tag{C.3}$$

which we use repeatedly below.

Step 1: Closed form for $\mathcal{V}'$ and non-negativity. Since $\mathbf{x}\mapsto f(\mathbf{x},\mathbf{B}^*)$ is a strictly convex quadratic with Hessian $\mathbf{M}+\mathbf{B}^*+\mathbf{C}$ and $\mathbf{x}^*$ is its unique minimizer, the standard quadratic expansion around $\mathbf{x}^*$ yields

$$f(\mathbf{x},\mathbf{B}^*)-f(\mathbf{x}^*,\mathbf{B}^*) \;=\; \frac{1}{2}\langle \mathbf{x}-\mathbf{x}^*,\,(\mathbf{M}+\mathbf{B}^*+\mathbf{C})(\mathbf{x}-\mathbf{x}^*)\rangle$$

for every $\mathbf{x}\in\mathbb{R}^n$. Setting $\mathbf{x}=\mathbf{x}^N$,

$$\mathcal{V}' \;=\; f(\mathbf{x}^N,\mathbf{B}^*)-f(\mathbf{x}^*,\mathbf{B}^*) \;=\; \frac{1}{2}\langle \mathbf{x}^N-\mathbf{x}^*,\,(\mathbf{M}+\mathbf{B}^*+\mathbf{C})(\mathbf{x}^N-\mathbf{x}^*)\rangle.$$

Under Property A, $\mathbf{M}\succ 0$, and $\mathbf{B}^*\succeq 0$, so $\mathbf{M}+\mathbf{B}^*+\mathbf{C}\succ 0$, hence $\mathcal{V}'\geq 0$.

Step 2: Closed form for $\mathcal{V}$. Using the naive first-order condition, $\langle \mathbf{x}^N,(\mathbf{M}+\mathbf{C})\mathbf{x}^N\rangle=\langle \psi^0+\psi,\mathbf{x}^N\rangle$, so

$$f(\mathbf{x}^N,\mathbf{B}^N) \;=\; -\frac{1}{2}\langle \psi^0+\psi,\mathbf{x}^N\rangle+\frac{1}{2}\langle \mathbf{x}^N,\mathbf{B}^N\mathbf{x}^N\rangle.$$

Similarly, the saddle-point FOC gives $\langle \mathbf{x}^*,(\mathbf{M}+\mathbf{B}^*+\mathbf{C})\mathbf{x}^*\rangle=\langle \psi^0+\psi,\mathbf{x}^*\rangle$, so

$$f(\mathbf{x}^*,\mathbf{B}^*) \;=\; -\frac{1}{2}\langle \psi^0+\psi,\mathbf{x}^*\rangle.$$

Subtracting,

$$\mathcal{V} \;=\; \frac{1}{2}\langle \mathbf{x}^N,\mathbf{B}^N\mathbf{x}^N\rangle-\frac{1}{2}\langle \psi^0+\psi,\mathbf{x}^N-\mathbf{x}^*\rangle. \tag{C.4}$$

We rewrite the second term using (C.3): by symmetry of $\mathbf{M}+\mathbf{C}$,

$$\langle \psi^0+\psi,\mathbf{x}^N-\mathbf{x}^*\rangle \;=\; \langle (\mathbf{M}+\mathbf{C})\mathbf{x}^N,\mathbf{x}^N-\mathbf{x}^*\rangle \;=\; \langle \mathbf{x}^N,(\mathbf{M}+\mathbf{C})(\mathbf{x}^N-\mathbf{x}^*)\rangle \;=\; \langle \mathbf{x}^N,\mathbf{B}^*\mathbf{x}^*\rangle.$$

Substituting into (C.4),

$$\mathcal{V} \;=\; \frac{1}{2}\langle \mathbf{x}^N,\mathbf{B}^N\mathbf{x}^N\rangle-\frac{1}{2}\langle \mathbf{x}^N,\mathbf{B}^*\mathbf{x}^*\rangle.$$

Step 3: $\mathcal{V}\geq\mathcal{V}'$. Cancelling $f(\mathbf{x}^*,\mathbf{B}^*)$,

$$\mathcal{V}-\mathcal{V}' \;=\; f(\mathbf{x}^N,\mathbf{B}^N)-f(\mathbf{x}^N,\mathbf{B}^*) \;=\; \frac{1}{2}\langle \mathbf{x}^N,(\mathbf{B}^N-\mathbf{B}^*)\mathbf{x}^N\rangle.$$

By Theorem 1-(b), $\mathbf{B}^N$ is the unique maximizer of $\mathbf{B}\mapsto\langle \mathbf{x}^N,\mathbf{B}\mathbf{x}^N\rangle$ over $\mathbf{B}\in\mathcal{B}$. Since

$\mathbf{B}^* \in \mathcal{B}$, the right-hand side is non-negative, hence $\mathcal{V} \geq \mathcal{V}'$.

Step 4: Equality conditions. If $\mathbf{v}_{ij} = 0$ for all $i, j \in N$, then $\mathcal{B} = \{\mathbf{0}\}$, so $\mathbf{B}^* = \mathbf{B}^N = \mathbf{0}$, $\mathbf{x}^N = \mathbf{x}^*$, and both $\mathcal{V}'$ and $\mathcal{V}$ vanish.

Conversely, suppose $\mathcal{V}' = 0$. Step 1 and $\mathbf{M} + \mathbf{B}^* + \mathbf{C} \succ 0$ together force $\mathbf{x}^N = \mathbf{x}^*$. Substituting into (C.3) gives $\mathbf{B}^*\mathbf{x}^* = \mathbf{0}$. Taking the inner product with $\mathbf{x}^*$ and expanding $\mathbf{B}^* = \sum_i \mathbf{q}_i(\mathbf{x}^*) \otimes \mathbf{q}_i(\mathbf{x}^*)$,

$$0 \;=\; \langle \mathbf{x}^*, \mathbf{B}^*\mathbf{x}^* \rangle \;=\; \sum_{i \in N} \langle \mathbf{q}_i(\mathbf{x}^*), \mathbf{x}^* \rangle^2 \;=\; \sum_{i \in N} \Big(\sum_{l \in N} \mathbf{v}_{il} |x_l^*| \Big)^2,$$

where the last equality uses $\langle \mathbf{q}_i(\mathbf{x}^*), \mathbf{x}^* \rangle = \sum_l s(x_l^*) \mathbf{v}_{il}\, x_l^* = \sum_l \mathbf{v}_{il} |x_l^*|$. Hence $\sum_l \mathbf{v}_{il} |x_l^*| = 0$ for all $i \in N$. By Property B, $|x_l^*| > 0$ for all l, so each non-negative summand $\mathbf{v}_{il}|x_l^*|$ must vanish, yielding $\mathbf{v}_{ij} = 0$ for all i, j.

For $\mathcal{V}$: from Step 3, $\mathcal{V} \geq \mathcal{V}' \geq 0$, so $\mathcal{V} = 0$ implies $\mathcal{V}' = 0$, and the conclusion follows from the previous case.

Step 5: Same-orthant case. If $\mathbf{x}^N$ and $\mathbf{x}^*$ lie in the same orthant, then $s(x_l^N) = s(x_l^*)$ for all $l \in N$, so $\mathbf{q}_i(\mathbf{x}^N) = \mathbf{q}_i(\mathbf{x}^*)$ for every i, hence $\mathbf{B}^N = \mathbf{B}^* =: \mathbf{B}$. Substituting into the expression from Step 2,

$$\mathcal{V} \;=\; \frac{1}{2}\langle \mathbf{x}^N, \mathbf{B}\mathbf{x}^N \rangle - \frac{1}{2}\langle \mathbf{x}^N, \mathbf{B}\mathbf{x}^* \rangle \;=\; \frac{1}{2}\langle \mathbf{x}^N, \mathbf{B}(\mathbf{x}^N - \mathbf{x}^*) \rangle.$$

■

Proof of Proposition 6. Recall that $\mathbf{x}^N = (\mathbf{M} + \mathbf{C})^{-1}(\psi^0 + \psi)$ does not depend on the variance profile, while

$$\mathbf{x}^*(\mathbf{v}) = \big(\mathbf{M} + \mathbf{C} + \mathbf{B}^*(\mathbf{v})\big)^{-1}(\psi^0 + \psi),$$
$$\mathbf{B}^*(\mathbf{v}) = \sum_{l=1}^{n} \mathbf{q}_l(\mathbf{x}^*) \otimes \mathbf{q}_l(\mathbf{x}^*),$$
$$\mathbf{B}^N(\mathbf{v}) = \sum_{l=1}^{n} \mathbf{q}_l(\mathbf{x}^N) \otimes \mathbf{q}_l(\mathbf{x}^N),$$

with $\mathbf{q}_l(\mathbf{x})_k = s(x_k)\mathbf{v}_{lk}$. Under Property B, $x_l^* \neq 0$ and $x_l^N \neq 0$ for all $l \in N$, so there exists a neighborhood of $\mathbf{v}$ on which $s(x_l^*)$ and $s(x_l^N)$ are constant for every l. On this neighborhood, $\mathbf{B}^*(\mathbf{v})$ is smooth in $\mathbf{v}$, and Property A gives $\mathbf{M} + \mathbf{B}^* + \mathbf{C} \succ 0$, so the implicit function theorem applied to the saddle-point first-order condition $(\mathbf{M} + \mathbf{B}^* + \mathbf{C})\mathbf{x}^* = \psi^0 + \psi$ implies that $\mathbf{x}^*$ is continuously differentiable in $\mathbf{v}$. We work in this neighborhood throughout, and define

$$\alpha_i^N = \sum_{l \in N} \mathbf{v}_{il} |x_l^N|, \quad \alpha_i^* = \sum_{l \in N} \mathbf{v}_{il} |x_l^*|, \quad \beta_i^N = \sum_{l \in N} s(x_l^*)\, \mathbf{v}_{il}\, x_l^N.$$

Step 1: Derivatives of $\mathbf{q}_l$. Since $s(x_k^N)$ and $s(x_k^*)$ are locally constant in $\mathbf{v}$, direct differentiation gives

$$\frac{\partial \mathbf{q}_l(\mathbf{x}^N)}{\partial \mathbf{v}_{ij}} = \delta_{li}\, s(x_j^N)\, \mathbf{e}_j, \qquad \frac{\partial \mathbf{q}_l(\mathbf{x}^*)}{\partial \mathbf{v}_{ij}}\bigg|_{\mathbf{x}^* \text{ fixed}} = \delta_{li}\, s(x_j^*)\, \mathbf{e}_j, \tag{C.5}$$

where $\mathbf{e}_j$ is the jth standard basis vector and δ_{li} is the Kronecker delta. Consequently,

$$\begin{aligned} \frac{\partial \mathbf{B}^N}{\partial \mathbf{v}_{ij}} &= s(x_j^N)\big[\mathbf{e}_j \otimes \mathbf{q}_i(\mathbf{x}^N) + \mathbf{q}_i(\mathbf{x}^N) \otimes \mathbf{e}_j\big], \\ \frac{\partial \mathbf{B}^*}{\partial \mathbf{v}_{ij}}\bigg|_{\mathbf{x}^* \text{ fixed}} &= s(x_j^*)\big[\mathbf{e}_j \otimes \mathbf{q}_i(\mathbf{x}^*) + \mathbf{q}_i(\mathbf{x}^*) \otimes \mathbf{e}_j\big]. \end{aligned}$$

Step 2: Derivative of $f(\mathbf{x}^*, \mathbf{B}^*)$. Recall $f(\mathbf{x}, \mathbf{B}) = \frac{1}{2}\langle \mathbf{x}, (\mathbf{M} + \mathbf{C})\mathbf{x}\rangle + \frac{1}{2}\langle \mathbf{x}, \mathbf{B}\mathbf{x}\rangle - \langle \psi^0 + \psi, \mathbf{x}\rangle$. By the chain rule,

$$\frac{d}{d\mathbf{v}_{ij}} f(\mathbf{x}^*, \mathbf{B}^*) = \nabla_{\mathbf{x}} f(\mathbf{x}^*, \mathbf{B}^*) \cdot \frac{\partial \mathbf{x}^*}{\partial \mathbf{v}_{ij}} + \frac{1}{2}\Big\langle \mathbf{x}^*, \frac{\partial \mathbf{B}^*}{\partial \mathbf{v}_{ij}}\bigg|_{\mathbf{x}^* \text{ fixed}} \mathbf{x}^*\Big\rangle.$$

The saddle-point first-order condition (5) yields $\nabla_{\mathbf{x}} f(\mathbf{x}^*, \mathbf{B}^*) = \mathbf{0}$, so the first term vanishes. Using Step 1 and the symmetry of $\partial \mathbf{B}^*/\partial \mathbf{v}_{ij}$,

$$\frac{d}{d\mathbf{v}_{ij}} f(\mathbf{x}^*, \mathbf{B}^*) = s(x_j^*)\, x_j^*\, \langle \mathbf{q}_i(\mathbf{x}^*), \mathbf{x}^*\rangle = |x_j^*|\, \alpha_i^*, \tag{C.6}$$

where $\langle \mathbf{q}_i(\mathbf{x}^*), \mathbf{x}^*\rangle = \sum_l s(x_l^*) \mathbf{v}_{il}\, x_l^* = \sum_l \mathbf{v}_{il} |x_l^*| = \alpha_i^*$.

Step 3: Derivative of $f(\mathbf{x}^N, \mathbf{B}^*)$ and equation (18). Since $\mathbf{x}^N$ does not depend on $\mathbf{v}$, only $\mathbf{B}^*$ contributes: $\frac{d}{d\mathbf{v}_{ij}} f(\mathbf{x}^N, \mathbf{B}^*) = \frac{1}{2}\Big\langle \mathbf{x}^N, \frac{d\mathbf{B}^*}{d\mathbf{v}_{ij}} \mathbf{x}^N\Big\rangle$. Since $\mathbf{q}_l$ depends on $\mathbf{x}^*$ only through the sign pattern, which is locally constant, the total derivative coincides with the partial in (C.5). Hence,

$$\frac{d}{d\mathbf{v}_{ij}} f(\mathbf{x}^N, \mathbf{B}^*) = s(x_j^*)\, x_j^N\, \langle \mathbf{q}_i(\mathbf{x}^*), \mathbf{x}^N\rangle = s(x_j^*)\, x_j^N\, \beta_i^N. \tag{C.7}$$

Combining (C.6) and (C.7),

$$\frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}} = \frac{d}{d\mathbf{v}_{ij}}\big[f(\mathbf{x}^N, \mathbf{B}^*) - f(\mathbf{x}^*, \mathbf{B}^*)\big] = s(x_j^*)\, x_j^N\, \beta_i^N - |x_j^*|\, \alpha_i^*,$$

which is (18).

Step 4: Derivative of $f(\mathbf{x}^N, \mathbf{B}^N)$ and equation (17). Since $\mathbf{x}^N$ does not depend on $\mathbf{v}$, an analogous computation using Step 1 yields

$$\frac{d}{d\mathbf{v}_{ij}} f(\mathbf{x}^N, \mathbf{B}^N) \;=\; s(x_j^N)\, x_j^N \,\langle \mathbf{q}_i(\mathbf{x}^N), \mathbf{x}^N \rangle \;=\; |x_j^N|\, \alpha_i^N. \tag{C.8}$$

Combining (C.6) and (C.8),

$$\frac{\partial \mathcal{V}}{\partial \mathbf{v}_{ij}} = \frac{d}{d\mathbf{v}_{ij}}\big[f(\mathbf{x}^N, \mathbf{B}^N) - f(\mathbf{x}^*, \mathbf{B}^*)\big] = |x_j^N|\, \alpha_i^N - |x_j^*|\, \alpha_i^*,$$

which is (17).

Step 5: Gap between costs and equation (19). Subtracting (18) from (17) and using $|x_j^N| = s(x_j^N)x_j^N$ and $x_l^N = s(x_l^N)|x_l^N|$,

$$\begin{aligned}\frac{\partial(\mathcal{V} - \mathcal{V}')}{\partial \mathbf{v}_{ij}} \;&=\; |x_j^N|\, \alpha_i^N - s(x_j^*)\, x_j^N\, \beta_i^N \;=\; x_j^N \sum_{l\in N} \mathbf{v}_{il}\, x_l^N \big[s(x_j^N)s(x_l^N) - s(x_j^*)s(x_l^*)\big] \\ &=\; |x_j^N| \sum_{l\in N} \mathbf{v}_{il}\, |x_l^N|\, s(x_j^N)s(x_l^N)\big[s(x_j^N)s(x_l^N) - s(x_j^*)s(x_l^*)\big].\end{aligned}$$

For each l, let $a_l = s(x_j^N)s(x_l^N) \in \{+1, -1\}$ and $b_l = s(x_j^*)s(x_l^*) \in \{+1, -1\}$. Since $a_l^2 = 1$, we have $a_l(a_l - b_l) = 1 - a_l b_l$, which equals 0 when $a_l = b_l$ (equivalently, $s(x_l^N)s(x_l^*) = s(x_j^N)s(x_j^*)$) and 2 otherwise. Hence,

$$\frac{\partial(\mathcal{V} - \mathcal{V}')}{\partial \mathbf{v}_{ij}} \;=\; 2|x_j^N| \sum_{\substack{l\in N \\ s(x_l^N)s(x_l^*) \neq s(x_j^N)s(x_j^*)}} \mathbf{v}_{il}\, |x_l^N| \;\geq\; 0,$$

which is (19).

Step 6: Same-orthant case and equation (20). Suppose $\mathbf{x}^N$ and $\mathbf{x}^*$ lie in the same orthant, so $s(x_l^N) = s(x_l^*)$ for all $l \in N$. Two consequences follow.

First, the summation index condition in (19) becomes empty, hence $\partial(\mathcal{V} - \mathcal{V}')/\partial \mathbf{v}_{ij} = 0$ and the two costs coincide.

Second, $\mathbf{q}_l(\mathbf{x}^N) = \mathbf{q}_l(\mathbf{x}^*)$ for all l when sign patterns agree, hence $\mathbf{B}^N = \mathbf{B}^*$. The same-orthant condition also gives $s(x_l^*)\, x_l^N = |x_l^N|$ and $s(x_j^*)\, x_j^N = |x_j^N|$, hence

$$\beta_i^N = \sum_{l\in N} s(x_l^*)\, \mathbf{v}_{il}\, x_l^N = \sum_{l\in N} \mathbf{v}_{il}\, |x_l^N| = \alpha_i^N.$$

Substituting into (18) yields

$$\frac{\partial \mathcal{V}'}{\partial \mathbf{v}_{ij}} \;=\; |x_j^N|\, \alpha_i^N - |x_j^*|\, \alpha_i^* \;=\; \frac{\partial \mathcal{V}}{\partial \mathbf{v}_{ij}},$$

which is (20). ∎

Online Appendix

OA Applications

OA.1 Network Expansion: Recruiting New Faculty

Illustration: $n = 2$ plus one entrant As an example, consider a network consisting of two existing members (agent 1 and agent 2) and one new agent (agent 3). For simplicity, we assume that the variance of the link is one for all agents; that is, all the diagonal entries of each $\overline{\mathbf{B}}_i$ are equal to one for $i = 1, 2, 3$. Consequently, each entry of $\overline{\mathbf{B}}_i$ matrix represents a correlation coefficient of the influences toward agent i. We assume that agent 1's influence on their own outcome is uncorrelated with agent 2's influence on agent 1's outcome (i.e., $(\mathbf{B}_1)_{12} = 0$). In contrast, we assume that agent 1's influence on agent 2's outcome is correlated with agent 2's influence on agent 2's outcome (i.e., $(\mathbf{B}_2)_{12} = \rho \in [-1, 1]$), and the exact correlation coefficient ρ is known to the DM. Finally, we suppose that all the other entries are unknown to the DM. Under these conditions, the following matrices illustrate the network uncertainty:

$$\overline{\mathbf{B}}_1 = \begin{array}{c} \\ 1 \\ 2 \\ 3 \end{array} \overset{\mathbf{B}_1 \;\; 1 \quad 2 \quad 3}{\begin{bmatrix} 1 & 0 & ? \\ 0 & 1 & ? \\ ? & ? & 1 \end{bmatrix}}, \quad \overline{\mathbf{B}}_2 = \begin{array}{c} \\ 1 \\ 2 \\ 3 \end{array} \overset{\mathbf{B}_2 \;\; 1 \quad 2 \quad 3}{\begin{bmatrix} 1 & \rho & ? \\ \rho & 1 & ? \\ ? & ? & 1 \end{bmatrix}}, \quad \text{and} \quad \overline{\mathbf{B}}_3 = \begin{array}{c} \\ 1 \\ 2 \\ 3 \end{array} \overset{1 \quad 2 \quad 3}{\begin{bmatrix} 1 & ? & ? \\ ? & 1 & ? \\ ? & ? & 1 \end{bmatrix}}.$$

There are two notable features in the examples of $\overline{\mathbf{B}}_3$ and the other two matrices. First, for any agent $i = 1, 2, 3$ in the network, all entries representing the correlation coefficient between the influence of the new agent 3 on i and the influence of another existing agent $j = 1, 2$ on i are undetermined; that is, $(\overline{\mathbf{B}}_i)_{j3}$ is undetermined. In the illustrated matrices, this feature is reflected by all the entries in the third row or column being undetermined and represented by the question mark (i.e., '?'). This highlights the first source of uncertainty introduced by the network expansion: the DM lacks information on how the new agent's influence on a given agent correlates with the influence exerted by other existing agents on that agent.

Second, the correlation coefficients of the existing members' influences on agents are asymmetric. Specifically, the influences of existing members on existing agents 1 and 2 are fully known. This is illustrated by the fact that, for agent $i = 1, 2$, the principal submatrix $\mathbf{B}_i$ of $\overline{\mathbf{B}}_i$, highlighted by a blue solid-line box in the illustrated matrices, has all its entries determined. In contrast, the influences of the existing agents on the new agent 3 are unknown. This highlights the second source of uncertainty introduced by the network expansion: the DM lacks information on how the existing agents' influences on the new agent are correlated with one another.

Consequently, all the off-diagonal entries in $\overline{\mathbf{B}}_3$ are undetermined and represented by the question mark in the illustrated matrices. This implies that, in the worst-case scenario, Nature's choice of $\overline{\mathbf{B}}_3$ presents the rank-1 structure, as established in Theorem 1. Therefore, the remaining analysis focuses on characterizing undetermined entries in $\overline{\mathbf{B}}_i$ for $i = 1, 2$ within an appropriate uncertainty set.

By linearity, we have $\langle \overline{\mathbf{x}}, \overline{\mathbf{B}}\overline{\mathbf{x}} \rangle = \langle \overline{\mathbf{x}}, \overline{\mathbf{B}}_1\overline{\mathbf{x}} \rangle + \langle \overline{\mathbf{x}}, \overline{\mathbf{B}}_2\overline{\mathbf{x}} \rangle + \langle \overline{\mathbf{x}}, \overline{\mathbf{B}}_3\overline{\mathbf{x}} \rangle$. Since none of the entries of $\overline{\mathbf{x}}$ is zero, Theorem 1 shows the unique $\overline{\mathbf{B}}_3^*$ is given by $\overline{\mathbf{B}}_3^* = \sigma_3^2(\mathbf{w}_3 \otimes \mathbf{w}_3)$ with $\sigma_3^2 = \sum_{j=1}^3 \mathbf{v}_{3j}^2 = 3$, $\mathbf{w}_3 = \frac{\mathbf{q}_3}{\|\mathbf{q}_3\|}$, and $\mathbf{q}_3 = (s(\overline{x}_1), s(\overline{x}_2), s(\overline{x}_3))^\intercal$. As a result, $\overline{\mathbf{B}}_3^*$ is a rank-1 matrix.

On the other hand, the uniqueness of the worst-case scenario for the existing agents $i = 1, 2$ stems from the strict convexity of the uncertainty set, which contrasts with the reasoning behind the uniqueness of $\overline{\mathbf{B}}_3$. To determine $\overline{\mathbf{B}}_1^*$, observe that $\langle \overline{\mathbf{x}}, \overline{\mathbf{B}}_1\overline{\mathbf{x}} \rangle = 2\overline{x}_3 \left((\overline{\mathbf{B}}_1)_{13}\overline{x}_1 + (\overline{\mathbf{B}}_1)_{23}\overline{x}_2\right) + \text{constant}$, where $(\overline{\mathbf{B}}_1)_{13} = \text{Corr}(\mathbf{G}_{11}, \mathbf{G}_{13}) \in [-1, 1]$ and $(\overline{\mathbf{B}}_1)_{23} = \text{Corr}(\mathbf{G}_{12}, \mathbf{G}_{13}) \in [-1, 1]$. The dashed square in Figure OA1-(a) represents these restrictions of the possible values of $(\overline{\mathbf{B}}_1)_{13}$ and $(\overline{\mathbf{B}}_1)_{23}$. Additionally, Nature's choice of uncertainty is constrained by the requirement that $\overline{\mathbf{B}}_1$ must be positive semi-definite, which holds if and only if $\det(\overline{\mathbf{B}}_1) \geq 0$ because $\mathbf{B}_1 \in \text{PD}^2$ and $(\overline{\mathbf{B}}_1)_{33} > 0$. The Schur Complement theorem (Meyer, 2010) states that $\det(\overline{\mathbf{B}}_1) \geq 0$ if and only if $(\overline{\mathbf{B}}_1)_{33} - ((\overline{\mathbf{B}}_1)_{13}, (\overline{\mathbf{B}}_1)_{23})\mathbf{B}_1^{-1}((\overline{\mathbf{B}}_1)_{13}, (\overline{\mathbf{B}}_1)_{23})^\intercal \geq 0$, which is equivalent to $(\overline{\mathbf{B}}_1)_{13}^2 + (\overline{\mathbf{B}}_1)_{23}^2 \leq 1$, represented by the gray unit disk in Figure OA1-(a) contained in the dashed square. Consequently, the uncertainty set is strictly convex and compact.[17] The worst-case scenario for agent 1 must therefore be chosen within this uncertainty set.

Since the uncertainty set is strictly convex, the gradient of $\langle \overline{\mathbf{x}}, \overline{\mathbf{B}}_1\overline{\mathbf{x}} \rangle$ with respect to $\left((\overline{\mathbf{B}}_1)_{13}, (\overline{\mathbf{B}}_1)_{23}\right)$ must be orthogonal to the boundary of the uncertainty set. The assumption of non-zero entries in $\overline{\mathbf{x}}$ implies that the gradient (represented by the

[17] Our proof in Appendix C addresses a more general condition, and the strict convexity of the uncertainty set is not a direct consequence of the Schur Complement theorem.

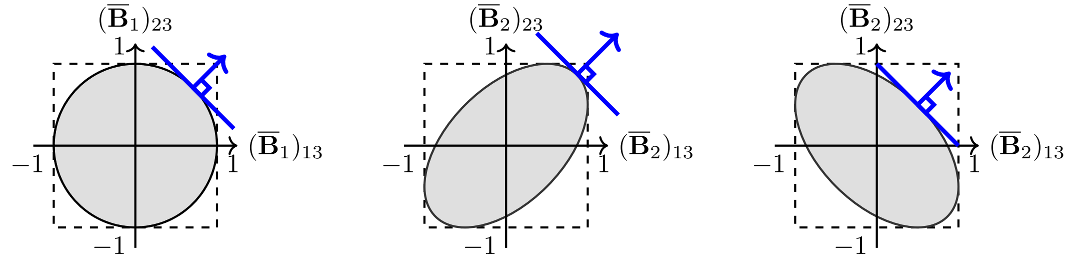

(a) uncertainty set for agent 1 (b) uncertainty set for agent 2 if $\rho = -0.5$ (c) uncertainty set for agent 2 if $\rho = 0.5$

Figure OA1: Illustration of the worst-case scenario. In each figure, the gradient of Nature's objective function at the unique worst-case scenario is orthogonal to the uncertainty set denoted by the gray elliptical disk.

blue arrow in Figure OA1) is not vanishing. Consequently, there is a unique optimal choice $\overline{\mathbf{B}}_1^*$, in which $\nabla\langle\overline{\mathbf{x}}, \overline{\mathbf{B}}_1\overline{\mathbf{x}}\rangle\big|_{\overline{\mathbf{B}}_1=\overline{\mathbf{B}}_1^*} = 2\overline{x}_3(\overline{x}_1, \overline{x}_2)$ is orthogonal to the boundary of the uncertainty set as illustrated in Figure OA1-(a). At the optimal choice, a trade-off arises among the correlation coefficients due to the strict convexity of the uncertainty set, which results from the partial correlation information constraint $\mathbf{B}_1$.

The worst-case scenario for agent 2, $\overline{\mathbf{B}}_2^*$, is also uniquely determined by the strict convexity of the uncertainty set. However, it is worth examining how the shape of uncertainty set changes due to the presence of the correlation among the influences toward agent 2. The dashed square in Figure OA1-(b) represents the constraints that $(\overline{\mathbf{B}}_2)_{13} = \mathrm{Corr}(\mathbf{G}_{21}, \mathbf{G}_{23}) \in [-1, 1]$ and $(\overline{\mathbf{B}}_2)_{23} = \mathrm{Corr}(\mathbf{G}_{22}, \mathbf{G}_{23}) \in [-1, 1]$. Again, Nature's choice of uncertainty faces another restriction: $\overline{\mathbf{B}}_2$ must be positive semi-definite, which holds if and only if, by the Schur Complement theorem,

$$\begin{bmatrix}(\overline{\mathbf{B}}_2)_{13} & (\overline{\mathbf{B}}_2)_{23}\end{bmatrix}\begin{bmatrix}1 & -\rho \\ -\rho & 1\end{bmatrix}\begin{bmatrix}(\overline{\mathbf{B}}_2)_{13} \\ (\overline{\mathbf{B}}_2)_{23}\end{bmatrix} \leq 1 - \rho^2.$$

The pairs of $\big((\overline{\mathbf{B}}_2)_{13}, (\overline{\mathbf{B}}_2)_{23}\big)$ satisfying the above inequality forms an elliptical disk as illustrated in Figure OA1-(b),(c). For each $\rho \in [-1, 1]$, there are two invariant principal components (i.e., eigenvectors) $(1, 1)^\mathsf{T}$ and $(1, -1)^\mathsf{T}$, with corresponding non-negative eigenvalues $1 - \rho$ and $1 + \rho$, respectively. The uncertainty set for $\rho = -0.5$ is shown as the rotated gray elliptical disk in Figure OA1-(b), and the uncertainty set for $\rho = 0.5$ is illustrated in Figure OA1-(c). In both figures, the uncertainty sets are strictly convex.

In summary, the strict convexity of the uncertainty sets yields a unique worst-case scenario $\overline{\mathbf{B}}_i^*$ for each $i = 1, 2, 3$, and moreover, $\overline{\mathbf{B}}_2^*$ can vary depending on the correlation coefficient ρ. If all the entries of $\overline{\mathbf{x}}$ are positive, the gradient of Nature's objective function at the worst-case scenario is orthogonal to the uncertainty set and lies in the first quadrant. As a result, the uncertainty set yields lower values of $(\overline{\mathbf{B}}_2^*)_{13}$ and $(\overline{\mathbf{B}}_2^*)_{23}$ as the correlation ρ increases, as shown in Figure OA1-(c). On the other hand, if $\overline{\mathbf{x}}$ includes a negative entry (e.g., $\overline{\mathbf{x}}_1 > 0$ and $\overline{\mathbf{x}}_2 < 0$), the worst-case scenario values for $(\overline{\mathbf{B}}_2)_{13}$ and $(\overline{\mathbf{B}}_2)_{23}$ may increase in magnitude as ρ increases, as in Figure OA1-(b).

OB Higher-Order Interactions

So far, we have addressed the robust intervention problem in the context of one-shot interactions among agents. However, in many studies in network economics, the focus shifts to settings where agents interact repeatedly or infinitely, leading to long-run equilibrium outcomes. Such scenarios are commonly analyzed using the equilibrium representation $(\mathbf{I} - \delta\mathbf{G})^{-1}$, where $\delta > 0$ captures the strength of the network effect.

If the spectral radius of $\delta\mathbf{G}$ is less than one, ensured when δ is sufficiently small, the equilibrium can be expressed as a power series: $(\mathbf{I} - \delta\mathbf{G})^{-1} = \sum_{k=0}^{\infty}(\delta\mathbf{G})^k$. In this representation, $\mathbf{G}_{ij}^k$ represents the weighted sum of all walks of length k from agent j to agent i, capturing the k-th order influence of agent j's allocation on agent i's outcome.[18] This power series representation aggregates the influence dynamics across all possible walks within the network.

The critical question, then, is how the DM can incorporate these higher-order influences and the associated higher-order uncertainties into the design of a robust intervention. By accounting for these dynamics, the DM can address the challenges of designing effective strategies in networks where agents' interactions propagate over time and across multiple pathways. Motivated by this, we now consider an extension of our model to incorporate higher-order uncertainty under additional assumptions.

First, we assume that with a second-order approximation of $(\mathbf{I} - \delta\mathbf{G})^{-1} \approx (\mathbf{I} + \delta\mathbf{G} + \delta^2\mathbf{G}^2)$, the DM minimizes the following objective:

$$\mathrm{E}\left[\|(\mathbf{I} + \delta\mathbf{G} + \delta^2\mathbf{G}^2)\mathbf{x} - \mathbf{z}\|^2\right]. \tag{OB.1}$$

Second, we assume $\mathbf{G}_{ii} = 0$ for all $i \in N$. This assumption is often valid in contexts such as network games or supply chain network models, where the zero-

[18] A walk of length k from node j to i is a sequence of nodes $(i_0, i_1, i_2, \ldots, i_k)$, such that $i_0 = j$, $i_k = i$, and nodes in the sequence need not be distinct (Jackson, 2010).

order interaction term $\mathbf{I}$ captures an agent's self-influence after normalization. Third, we assume statistical independence between the rows of $\mathbf{G}$, meaning that the influence on an agent is independent of the influence on any other agents. For instance, in a social learning model, agent i's weight on agent k is independent of agent j's weight on k.

Under these assumptions, the adversarial Nature's choice of the worst-case scenario is characterized by a unique rank-1 covariance matrix $\mathbf{B}_i^*$ for each agent i. To simplify exposition, we assume that $\mathbf{m}_{ij} = \mathrm{E}[\mathbf{G}_{ij}] = 0$ for all $i, j \in N$, allowing us to focus solely on the impact of uncertainty.[19] Under this assumption, the expected squared distance between agent i's final outcome and the target outcome z_i is

$$\mathrm{E}\left[|(\mathbf{I} + \delta\mathbf{G} + \delta^2\mathbf{G}^2)_i\mathbf{x} - z_i|^2\right] = x_i^2 + \delta^2\langle\mathbf{x}, \mathbf{B}_i\mathbf{x}\rangle + \sum_{k=1}^{n} \delta^4\mathbf{v}_{ik}^2\langle\mathbf{x}, \mathbf{B}_k\mathbf{x}\rangle + z_i^2 - 2z_ix_i. \tag{OB.2}$$

The term $\delta^2\langle\mathbf{x}, \mathbf{B}_i\mathbf{x}\rangle$ in expression OB.2 arises from $(\delta\mathbf{G})_i\mathbf{x}$, capturing the uncertainty generated by the first-order interactions among the agents with respect to agent i.

Another term, $\delta^4\sum_{k=1}^{n}\mathbf{v}_{ik}^2\langle\mathbf{x}, \mathbf{B}_k\mathbf{x}\rangle$, accounts for the second-order interactions affecting agent i's final outcome. For these second-order interactions, consider pairs of agents s and k with intermediate agents s' and k' directly influencing agent i. The second-order interactions from s to i and from k to i are represented by $\mathbf{G}_{is'}\mathbf{G}_{s's}$ and $\mathbf{G}_{ik'}\mathbf{G}_{k'k}$, respectively. If $s' \neq k'$, the correlation between these second-order interactions is zero due to the independence in link formation among the agents' interactions, captured by the computation $\mathrm{E}\left[\mathbf{G}_{is'}\mathbf{G}_{s's}\mathbf{G}_{ik'}\mathbf{G}_{k'k}\right] = \mathrm{E}\left[\mathbf{G}_{is'}\mathbf{G}_{ik'}\right]\mathrm{E}\left[\mathbf{G}_{s's}\right]\mathrm{E}\left[\mathbf{G}_{k'k}\right] = 0$.

Consequently, the second order effect arises only when there is a common intermediate agent (i.e., $k' = s'$). For each intermediate agent, the correlations among walks of length 2 that share agent k' as the common intermediate agent are amplified by the discounted variance $\delta^4\mathbf{v}_{ik'}^2$. For example, in Figure OA2, the solid blue arrows directed toward intermediate agent 2 influence agent 1's outcome and are independent of the other arrows directed toward a distinct agent, such as the dashed green arrows leading to intermediate agent 3. Then, the impact of the blue arrows on agent 1 is weighted by the discounted variance term $\delta^4\mathbf{v}_{12}^2$, while the impact of the green arrows on agent 1 is weighted by the discounted variance term $\delta^4\mathbf{v}_{13}^2$.

[19] The proof of Proposition A4 does not rely on this zero mean influence of the agents.

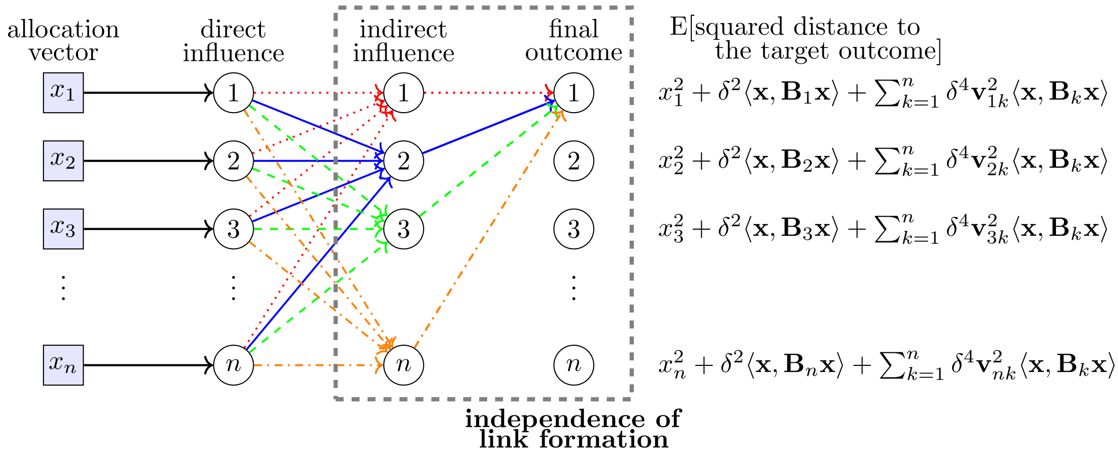

Figure OA2: Illustration of second-order uncertainty generation in network. For simplicity, here we assume that $\mathbf{z} = \mathbf{0}$.

By summing the squared distances for all agents, we obtain

$$\mathrm{E}\left[\|(\mathbf{I} + \delta \mathbf{G} + \delta^2 \mathbf{G}^2)\mathbf{x} - \mathbf{z}\|^2\right] = \langle \mathbf{x}, \mathbf{x} \rangle + \sum_{i=1}^{n} w_i \langle \mathbf{x}, \mathbf{B}_i \mathbf{x} \rangle - 2\langle \mathbf{z}, \mathbf{x} \rangle + \langle \mathbf{z}, \mathbf{z} \rangle, \qquad \text{OB.3}$$

where $w_i = \delta^2(1 + \delta^2 \sum_{k=1}^n \mathbf{v}_{ki}^2) > 0$. The above expression remains linear in each $\mathbf{B}_i$, similar to expression (3) in the main model without higher-order considerations. As detailed in Theorem 1, this implies that the adversarial Nature's selection of the worst-case scenario satisfies the rank-1 property for all agents. The following proposition summarizes this result.

Proposition A4 *With the higher-order consideration with the objective function OB.2, there exists a unique worst-case scenario $\mathbf{B}_i^*$ for each agent $i \in N$, and it is rank-1.*

The uniqueness of the worst-case scenario, in turn, allows the DM to solve her robust optimal intervention problem by selecting her intervention based on the first-order condition under the worst-case scenario, in a manner analogous to Theorem 1.

We conclude this subsection by noting that for k-th order interaction considerations with $k \geq 3$, the linearity in the objective with respect to the uncertainties $(\mathbf{B}_i)_{i \in N}$ no longer holds, even under the assumption of independent link formation. The reason is that, even with independent link formation, higher-order interactions among influences on an agent can emerge. For example, consider walks of length 3, such as $(3, 2, 1, 2)$ and $(4, 2, 1, 2)$, which differ only in the initial agent. In this case, the interaction among the influences $\mathbf{G}_{23}$, $\mathbf{G}_{24}$ and $\mathbf{G}_{21}$ in the covariance term may introduce nonlinear effects. As a result, the rank-1 property and the uniqueness of the worst-case scenario may no longer hold. We leave the analysis of higher-order interactions and robust intervention design for future research.

Proof of Proposition A4

Proof. For simplicity, we prove the proposition by assuming that the mean influence of the link $\mathbf{G}_{ij}$ is zero for all $i, j \in N$. Then, we provide a general proof without the assumption.

For each $i \in N$, we find that

$$|(\mathbf{I} + \delta\mathbf{G} + \delta^2\mathbf{G}^2)_i\mathbf{x} - z_i|^2 = \left|x_i + \delta\sum_{j=1}^{n}\mathbf{G}_{ij}x_j + \delta^2\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l - z_i\right|^2$$

$$= x_i^2 + \left(\delta\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right)^2 + \left(\delta^2\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)^2 + z_i^2 + 2x_i\left(\delta\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right)$$
$$+ 2x_i\left(\delta^2\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right) - 2x_iz_i + 2\left(\delta\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right)\left(\delta^2\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)$$
$$- 2z_i\left(\delta\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right) - 2z_i\left(\delta^2\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right). \quad \text{OB.4}$$

We investigate the expectation of each term in expression OB.4. First, we find that

$$\mathrm{E}\left[\left(\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)\right] = 0, \quad \mathrm{E}\left[\left(\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right)\left(\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)\right] = 0,$$
$$\mathrm{E}\left[\left(\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right)\right] = 0, \quad \mathrm{E}\left[\left(\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)\right] = 0.$$

Second, $\mathrm{E}\left[\left(\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right)^2\right] = \langle\mathbf{x}, \mathbf{B}_i\mathbf{x}\rangle$ as in the main model. Third, we observe that

$$\mathrm{E}\left[\left(\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)^2\right] = \mathrm{E}\left[\left(\sum_{1\le k,l,s,t\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}\mathbf{G}_{is}\mathbf{G}_{st}x_lx_t\right)\right]$$

$$= \sum_{k\ne i,s\ne i,l,t}\mathrm{E}\left[\mathbf{G}_{ik}\mathbf{G}_{is}\right]\mathrm{E}\left[\mathbf{G}_{kl}\mathbf{G}_{st}\right]x_lx_t = \sum_{k\ne i,l,t}\mathrm{E}\left[\mathbf{G}_{ik}^2\right]\mathrm{E}\left[\mathbf{G}_{kl}\mathbf{G}_{kt}\right]x_lx_t = \sum_{k=1}^{n}\mathbf{v}_{ik}^2\langle\mathbf{x}, \mathbf{B}_k\mathbf{x}\rangle.$$

Thus, we obtain expression OB.2 in the main text:

$$\mathrm{E}\left[|(\mathbf{I}+\delta\mathbf{G}+\delta^2\mathbf{G}^2)_i\mathbf{x}-z_i|^2\right] = x_i^2 + \mathrm{E}\left[\left(\delta\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right)^2 + \left(\delta^2\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)^2\right] + z_i^2$$

$$= x_i^2 + \delta^2\langle\mathbf{x},\mathbf{B}_i\mathbf{x}\rangle + \delta^4\sum_{k=1}^{n}\mathbf{v}_{ik}^2\langle\mathbf{x},\mathbf{B}_k\mathbf{x}\rangle + z_i^2 - 2z_ix_i. \qquad \text{OB.5}$$

Finally, by summing the squared distances for all agents, we obtain

$$\mathrm{E}\left[\|(\mathbf{I}+\mathbf{G}+\mathbf{G}^2)\mathbf{x}-\mathbf{z}\|^2\right] = \sum_{i=1}^{n}\left(x_i^2 + \delta^2\langle\mathbf{x},\mathbf{B}_i\mathbf{x}\rangle + \delta^4\sum_{k=1}^{n}\mathbf{v}_{ik}^2\langle\mathbf{x},\mathbf{B}_k\mathbf{x}\rangle - 2z_ix_i + z_i^2\right)$$

$$= \langle\mathbf{x},\mathbf{x}\rangle + \sum_{i=1}^{n}w_i\langle\mathbf{x},\mathbf{B}_i\mathbf{x}\rangle - 2\langle\mathbf{z},\mathbf{x}\rangle + \langle\mathbf{z},\mathbf{z}\rangle, \qquad \text{OB.6}$$

where $w_i = \delta^2 + \delta^4\sum_{k=1}^{n}\mathbf{v}_{ki}^2 > 0$ as in expression OB.3 in the main text.

Since the DM's objective function OB.6 is linear in each $\mathbf{B}_i$, the rank-1 property and the uniqueness of $\mathbf{B}_i^*$ are held by Theorem 1. Therefore, the worst-case scenario $\mathbf{B}^* = \sum_{i=1}^{n}\mathbf{B}_i^*$ is unique.

We now prove the proposition without the assumption. To economize notation, without loss of generality, we let $\delta = 1$ because it does not affect the linearity of the DM's objective function. We let $\overline{\mathbf{G}} = \mathrm{E}[\mathbf{G}]$ and write $\mathbf{G} = \overline{\mathbf{G}} + \mathbf{U}$. We denote by $\mathbf{m}_i$ and $\mathbf{m}^i$ the ith row vector and column vector of $\overline{\mathbf{G}}$, respectively. Recall that $\mathbf{M}_i = \mathbf{m}_i \otimes \mathbf{m}_i$, and $\mathrm{E}[\mathbf{U}_i\mathbf{U}_i^\intercal] = \mathbf{B}_i$. Since we assume that $\mathbf{G}_{ii} = 0$, we have $\mathbf{U}_{ii} = \mathbf{m}_{ii} = 0$. We calculate that for each $i \in N$,

$$\mathrm{E}\left[\left(\sum_{j=1}^{n}\mathbf{G}_{ij}x_j\right)^2\right] = \mathrm{E}\left[\left((\mathbf{m}_i+\mathbf{U}_i)^\intercal\mathbf{x}\right)^2\right] = \langle\mathbf{x},\mathbf{B}_i\mathbf{x}\rangle + \langle\mathbf{x},\mathbf{M}_i\mathbf{x}\rangle.$$

In addition, it follows that for each $i \in N$,

$$\mathrm{E}\left[\left(\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)^2\right] = \mathrm{E}\left[\sum_{1\le k,l,s,t\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}\mathbf{G}_{is}\mathbf{G}_{st}x_lx_t\right]$$

$$= \sum_{1\le l,t\le n}\sum_{k\ne i,s\ne i}\mathrm{E}\left[\mathbf{G}_{ik}\mathbf{G}_{is}\right]\mathrm{E}\left[\mathbf{G}_{kl}\mathbf{G}_{st}\right]x_lx_t$$

$$= \sum_{1\le l,t\le n}\sum_{k\ne i,s\ne i}(\mathrm{E}[\mathbf{U}_{ik}\mathbf{U}_{is}] + \mathbf{m}_{ik}\mathbf{m}_{is})\,\mathrm{E}([\mathbf{U}_{kl}\mathbf{U}_{st}] + \mathbf{m}_{kl}\mathbf{m}_{st})x_lx_t.$$

We find the following useful expressions for each $i \in N$:

$$\sum_{1\le l,t\le n}\sum_{k\neq i,s\neq i}\mathrm{E}[\mathbf{U}_{ik}\mathbf{U}_{is}]\,\mathrm{E}\left[\mathbf{U}_{kl}\mathbf{U}_{st}\right]x_l x_t = \sum_{k\neq i}\sum_{l,t}\mathrm{E}[\mathbf{U}_{ik}\mathbf{U}_{ik}]\,\mathrm{E}\left[\mathbf{U}_{kl}\mathbf{U}_{kt}\right]x_l x_t = \sum_{k}\mathbf{v}_{ik}^2\langle\mathbf{x},\mathbf{B}_k\mathbf{x}\rangle,$$

$$\sum_{1\le l,t\le n}\sum_{k\neq i,s\neq i}\mathbf{m}_{ik}\mathbf{m}_{is}\,\mathrm{E}\left[\mathbf{U}_{kl}\mathbf{U}_{st}\right]x_l x_t = \sum_{k}\mathbf{m}_{ik}^2\langle\mathbf{x},\mathbf{B}_k\mathbf{x}\rangle,$$

$$\sum_{1\le l,t\le n}\sum_{k\neq i,s\neq i}\mathbf{m}_{kl}\mathbf{m}_{st}\,\mathrm{E}[\mathbf{U}_{ik}\mathbf{U}_{is}]x_l x_t = \sum_{l,t}\langle x_l\mathbf{m}^l,\mathbf{B}_i(x_t\mathbf{m}^\intercal)\rangle = \langle\overline{\mathbf{G}}\mathbf{x},\mathbf{B}_i\overline{\mathbf{G}}\mathbf{x}\rangle,$$

$$\sum_{1\le l,t,k,s\le n}\mathbf{m}_{ik}\mathbf{m}_{is}\mathbf{m}_{kl}\mathbf{m}_{st}x_l x_t = \sum_{l,t}(\overline{\mathbf{G}}_{il}^2 x_l)(\overline{\mathbf{G}}_{it}^2 x_t) = (\overline{\mathbf{G}}^2\mathbf{x})_i^2.$$

Then, we obtain that for each $i \in N$,

$$\mathrm{E}\left[\left(\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right)^2\right] = \sum_{k=1}^n(\mathbf{v}_{ik}^2+\mathbf{m}_{ik}^2)\langle\mathbf{x},\mathbf{B}_k\mathbf{x}\rangle + \langle\overline{\mathbf{G}}\mathbf{x},\mathbf{B}_i\overline{\mathbf{G}}\mathbf{x}\rangle + (\overline{\mathbf{G}}^2\mathbf{x})_i^2,$$

$$\mathrm{E}\left[\sum_{j=1}^n\mathbf{G}_{ij}x_j\right] = \sum_{j=1}^n\mathbf{m}_{ij}x_j = \langle\mathbf{m}_i,\mathbf{x}\rangle = (\overline{\mathbf{G}}\mathbf{x})_i,$$

$$\mathrm{E}\left[\sum_{1\le k,l\le n}\mathbf{G}_{ik}\mathbf{G}_{kl}x_l\right] = \sum_{k,l}\left(\mathbf{m}_{ik}\mathbf{m}_{kl}+\mathrm{E}\left[\mathbf{U}_{ik}\mathbf{U}_{kl}\right]\right)x_l = (\overline{\mathbf{G}}^2\mathbf{x})_i + 0 = (\overline{\mathbf{G}}^2\mathbf{x})_i,$$

$$\begin{aligned}\mathrm{E}\left[\sum_{1\le j,k,l\le n}\mathbf{G}_{ij}\mathbf{G}_{ik}\mathbf{G}_{kl}x_j x_l\right] &= \sum_{1\le j,l\le n}\sum_{k\neq i}\mathrm{E}\left[\mathbf{G}_{ij}\mathbf{G}_{ik}\right]\mathrm{E}\left[\mathbf{G}_{kl}\right]x_j x_l\\ &= \sum_{1\le j,l,k\le n}\left(\mathrm{E}\left[\mathbf{U}_{ij}\mathbf{U}_{ik}\right]+\mathbf{m}_{ij}\mathbf{m}_{ik}\right)\mathbf{m}_{kl}x_j x_l\\ &= \sum_{l=1}^n\left(\langle x_l\mathbf{m}^l,\mathbf{B}_i\mathbf{x}\rangle + \langle x_l\mathbf{m}^l,\mathbf{M}_i\mathbf{x}\rangle\right)\\ &= \langle\overline{\mathbf{G}}\mathbf{x},\mathbf{B}_i\mathbf{x}\rangle + \langle\overline{\mathbf{G}}\mathbf{x},\mathbf{M}_i\mathbf{x}\rangle.\end{aligned}$$

Thus, we combine the above expressions and obtain that for each $i \in N$,

$$\begin{aligned}&\mathrm{E}\left[|(\mathbf{I}+\mathbf{G}+\mathbf{G}^2)_i\mathbf{x}-z_i|^2\right]\\ &= \langle\mathbf{x},\mathbf{B}_i\mathbf{x}\rangle + \sum_{k=1}^n(\mathbf{v}_{ik}^2+\mathbf{m}_{ik}^2)\langle\mathbf{x},\mathbf{B}_k\mathbf{x}\rangle + \langle\overline{\mathbf{G}}\mathbf{x},\mathbf{B}_i\overline{\mathbf{G}}\mathbf{x}\rangle + 2\langle\overline{\mathbf{G}}\mathbf{x},\mathbf{B}_i\mathbf{x}\rangle + \langle\mathbf{x},\mathbf{M}_i\mathbf{x}\rangle + (\overline{\mathbf{G}}^2\mathbf{x})_i^2\\ &\quad + 2\langle\overline{\mathbf{G}}\mathbf{x},\mathbf{M}_i\mathbf{x}\rangle + 2\langle x_i\mathbf{m}_i,\mathbf{x}\rangle + 2x_i(\overline{\mathbf{G}}^2\mathbf{x})_i + x_i^2 - 2\langle z_i\mathbf{m}_i,\mathbf{x}\rangle - 2z_i(\overline{\mathbf{G}}^2\mathbf{x})_i - 2z_i x_i + z_i^2.\end{aligned}$$

We now let $\alpha_k = \sum_{i=1}^n(\mathbf{v}_{ik}^2+\mathbf{m}_{ik}^2)$. By summing over i, it follows that

$$\mathrm{E}\left[\|(\mathbf{I}+\mathbf{G}+\mathbf{G}^2)\mathbf{x}-\mathbf{z}\|^2\right] = \sum_{i=1}^n\left(\alpha_i\langle\mathbf{x},\mathbf{B}_i\mathbf{x}\rangle + \langle(\overline{\mathbf{G}}+\mathbf{I})\mathbf{x},\mathbf{B}_i(\overline{\mathbf{G}}+\mathbf{I})\mathbf{x}\rangle\right) + Q(\mathbf{x}), \quad \text{OB.7}$$

where $Q(\mathbf{x})$ is a quadratic function of $\mathbf{x}$ not involving the uncertainty matrix $(\mathbf{B}_i)_i$ for any $i \in N$. As such, we can consider Q as an additional cost part of the DM's objective function. It is straightforward to see that the DM's objective function OB.7 is still a linear function of $(\mathbf{B}_i)_i$ as before. Consequently, we conclude that when $(\mathbf{B}_1^*, \ldots, \mathbf{B}_n^*)$ is Nature's best response with respect to a given $\mathbf{x}$ having no zero entry, its ith component $\mathbf{B}_i^*$ is uniquely determined as a rank-1 matrix if and only if the following nondegeneracy condition holds:

$$\alpha_i x_j x_k + \big((\overline{\mathbf{G}} + \mathbf{I})\mathbf{x}\big)_j \big((\overline{\mathbf{G}} + \mathbf{I})\mathbf{x}\big)_k \neq 0 \quad \text{for all } 1 \leq j < k \leq n. \qquad \text{OB.8}$$

Therefore, the proposition is proven. ■

OC Microfoundation of the Uncertainty Set

This appendix collects additional interpretations of the uncertainty set $\mathcal{B}$ introduced in Subsection 2.2. Throughout, $\mathbf{U}_i = \mathbf{G}_i^\intercal - \mathbf{m}_i \in \mathbb{R}^n$ denotes the column vector of row-i deviations, as defined in Subsection 2.1.

OC.1 Bayesian Interpretation: A Set of Priors Consistent with Trusted Moments

Let Π be the set of distributions over $\mathbf{G}$ such that $\mathrm{E}_\pi[\mathbf{G}_{ij}] = \mathbf{m}_{ij}$ and $\mathrm{Var}_\pi(\mathbf{G}_{ij}) = \mathbf{v}_{ij}^2$ for all $i, j \in N$. For any $\pi \in \Pi$, define $\mathbf{B}_i(\pi) = \mathrm{E}_\pi[\mathbf{U}_i \mathbf{U}_i^\intercal]$. The mapping $\pi \mapsto (\mathbf{B}_i(\pi))_{i \in N}$ induces a set of feasible row covariance matrices contained in $\prod_{i \in N} \mathcal{B}_i$, and Nature's move can be interpreted as selecting a least favorable dependence structure among those consistent with the DM's trusted marginal information. Moreover, $\mathcal{B}_i$ can be viewed as a least-restrictive envelope: for any $\mathbf{B}_i \succeq 0$ with diagonal $(\mathbf{v}_{i1}^2, \ldots, \mathbf{v}_{in}^2)$, there exists a distribution over $\mathbf{U}_i$ with covariance $\mathbf{B}_i$ (for example, multivariate normal), so $\mathcal{B}_i$ contains all covariance structures compatible with the maintained second-moment information and covariance feasibility.

OC.2 Partial-Identification Interpretation: An Identified Set for Dependence

Suppose an econometrician observes data generated by $\mathbf{G}$ and can consistently estimate $\mathbf{m}_{ij}$ and $\mathbf{v}_{ij}^2$ for each link, but cannot identify $\mathrm{Cov}(\mathbf{G}_{ij}, \mathbf{G}_{ik})$ beyond general inequalities implied by feasibility (e.g., Cauchy–Schwarz and positive semidefiniteness). Then $\mathcal{B}_i$ is an identified set of observationally equivalent within-row covariance structures, and $\mathcal{B}$ is the corresponding identified set for aggregated covariance objects that matter for expected loss. In this interpretation, the DM's max–min problem is a minimax decision rule over the identified set, as in robust treatment-choice problems under ambiguity and partial identification.[20]

OC.3 Data-Based Microfoundations

Microfoundation 1: link-specific measurement identifies marginals but not correlations. In many network applications, different links (i, j) are measured using distinct data sources, instruments, or experimental designs, so the DM can estimate each $\mathbf{m}_{ij}$ and $\mathbf{v}_{ij}^2$ from link-by-link variation but does not observe joint realizations that would identify covariances across (i, j) and (i, k) for fixed i and distinct j, k. The resulting statistical information disciplines the diagonal of $\mathbf{B}_i$ but leaves the off-diagonals largely unrestricted, motivating the use of $\mathcal{B}_i$.

Microfoundation 2: limited overlap in panels or rotating samples. Even when all links are conceptually observed in a single dataset, limitations such as short panels, rotating samples, or missing-by-design observations may imply that, for each (i, j), there is enough repeated variation to estimate $\mathrm{Var}(\mathbf{G}_{ij})$ but insufficient overlap to estimate $\mathrm{Cov}(\mathbf{G}_{ij}, \mathbf{G}_{ik})$ with precision. In such settings, treating off-diagonal dependence as ambiguous while trusting the marginal second moments is a disciplined reduced-form approximation, and worst-case evaluation over $\mathcal{B}_i$ corresponds to a conservative policy criterion that is valid across dependence patterns consistent with the data.

Microfoundation 3: latent common shocks generate comovement that is hard to discipline. Flexible structural sources of dependence are latent common shocks affecting multiple links toward the same receiver i. For instance, represent row-i deviations as $(\mathbf{U}_i)_j = a_{ij} f_i + u_{ij}$ for $j \in N$, where f_i is a receiver-i common shock with $\mathrm{Var}(f_i) = \sigma_i^2$, a_{ij} are link-specific loadings, and u_{ij} is idiosyncratic noise with $\mathrm{E}[u_{ij}] = 0$ and $\mathrm{Var}(u_{ij}) = \mathbf{v}_{ij}^2 - a_{ij}^2 \sigma_i^2$, with $a_{ij}^2 \sigma_i^2 \leq \mathbf{v}_{ij}^2$. Many combinations of

[20] For related approaches to robust decision making with partially identified models, see, for example, Stoye (2012); Montiel Olea et al. (2025).

$(a_{ij})_{j \in N}$ and σ_i^2 are consistent with the same marginal variances $\mathbf{v}_{ij}^2$ but imply different covariance patterns $\mathrm{Cov}(\mathbf{G}_{ij}, \mathbf{G}_{ik}) = a_{ij} a_{ik} \sigma_i^2$ across (j, k). When the DM can estimate $\mathbf{v}_{ij}^2$ but cannot credibly bound the strength of the common component, $\mathcal{B}_i$ captures the implied ambiguity without committing to a particular factor specification.

OC.4 Why Receiver-Specific Ambiguity and Aggregation are Natural

The object $\mathbf{U}_i$ aggregates the uncertainty relevant for receiver i because it determines the random component of $(\mathbf{Gx})_i$. Accordingly, the maintained information and the ambiguity are naturally formulated row-by-row: the DM's knowledge concerns marginal moments of each $\mathbf{G}_{ij}$, while the economically relevant unknowns are the covariances among links entering a common receiver. The aggregate set $\mathcal{B}$ arises from combining receiver-specific covariance choices $\mathbf{B}_i \in \mathcal{B}_i$ via $\mathbf{B} = \sum_{i=1}^{n} \mathbf{B}_i$, yielding a transparent interpretation of Nature's move: Nature selects a feasible dependence structure within each row, subject to the DM's trusted marginal second moments and covariance feasibility, to maximize the DM's expected loss under the chosen intervention.