TARGETING INTERVENTIONS IN NETWORKS

ANDREA GALEOTTI, BENJAMIN GOLUB, AND SANJEEV GOYAL

ABSTRACT. Individuals interact strategically with their network neighbors. A planner can shape incentives in pursuit of an aggregate goal, such as maximizing welfare or minimizing volatility. We analyze a variety of targeting problems by identifying how a given profile of incentive changes is amplified or attenuated by the strategic spillovers in the network. The optimal policies are simplest when the budget for intervention is large. If actions are strategic complements, the optimal intervention changes all agents' incentives in the same direction and does so in proportion to their eigenvector centralities. In games of strategic substitutes, the optimal intervention is very different: it moves neighbors' incentives in opposite directions, dividing local communities into positively and negatively targeted agents, with few links across these two categories. To derive these results and characterize optimal interventions more generally, we introduce a method of decomposing any potential intervention into *principal components* determined by the network. A particular ordering of principal components describes the planner's priorities across a range of network intervention problems.

Date Printed. March 12, 2018.
We are grateful for conversations with Xavier Vives and Glen Weyl at an earlier stage of this project. Gustavo Nicolas Paez and Eduard Talamás provided exceptional research assistance. We thank Sihua Ding, Joerg Kalbfuss, Fakhteh Saadatniaki, Alan Walsh, and Yves Zenou for detailed readings of earlier drafts. Andrea Galeotti is grateful to the European Research Council for support through the ERC-consolidator grant (award no. 724356) and to the European University Institute for support through the Internal Research Grant.
Galeotti: Department of Economics, European University Institute, and Department of Economics, University of Essex, Andrea.Galeotti@eui.eu. Golub: Department of Economics, Harvard University and University of Pennsylvania, bgolub@fas.harvard.edu; Goyal: Faculty of Economics and Christ's College, University of Cambridge, sg472@cam.ac.uk.

1. Introduction

Consider a group of individuals who interact strategically, and an external entity—a *planner*—who seeks to achieve a goal through an intervention that changes individuals' incentives. For concreteness, we present three applications.

(1) School pupils choose their level of educational effort. A pupil's incentives to study are affected by his own attributes (ability, etc.) and the efforts exerted by his friends. A school principal seeks to improve total student welfare and the instruments at her disposal include offering additional tutoring or rewards for achievement to some students.
(2) Firms invest in capacity, optimizing in response to costs and benefits idiosyncratic to their own business, as well as other firms' investment decisions. A government seeks to reduce the volatility of aggregate investment and can affect the variances and covariances of firms' idiosyncratic incentives for investment—e.g., by stabilizing prices of their inputs.
(3) Firms in a supply chain set prices; these decisions affect the demand, and hence pricing, of other firms. A regulator seeks to maximize total market surplus. As in (2), it can, at a cost, affect the variability and correlations of the prices of firms' inputs.

These cases can be studied within a common framework of interventions in a network game. Agents' payoffs depend on their own attributes, their actions, and the actions of some others, whom we call their contacts or neighbors. Actions may be strategic complements or strategic substitutes. A planner can, by allocating limited intervention resources, change the incentives of some targeted individuals. Through the network, this will also change the incentives of their neighbors. The planner seeks to improve some aggregate outcome such as total welfare, volatility, or market surplus. We study the question: how should interventions be targeted?

To describe some of the basic forces, let us start by considering setting (1)—e.g., pupils studying in school—under strategic complements. Suppose that the planner can target a given individual and encourage him to work harder. His raised effort, in turn, pushes up the efforts of his friends due to the strategic complementarity, which, in turn, increases the efforts of their friends, and so forth. It is natural to suppose that the planner's marginal costs of raising any one individual's effort are increasing. Thus, given a budget, the planner may be best off targeting multiple individuals, rather than focusing all effort on one. But how should she allocate the effort?

An important observation is that in the strategic complements case, the planner will want to move neighbors' incentives together, since increasing someone's effort makes it easier to increase the efforts of his neighbors. In other words, the planner's investments to encourage two neighbors to work harder are complementary. The optimal policy will, then, exploit this complementarity to amplify the impact of interventions. The first main result says that, under certain conditions, the best amplification is achieved by targeting agents in proportion to their *eigenvector centrality*, which measures how well-connected they are in the network as a whole. An agent's eigenvector centrality is the right measure of the total direct and indirect complementarity between changes to this agent's incentives and those of everyone else.

Next, consider the case of strategic substitutes. To take a concrete example, consider a decentralized team collaborating to develop a software product. Effort by one individual towards documenting and testing code reduces the incentives of others using the same code (his network neighbors) to do the same. A planner can encourage some team members to invest more in this activity, e.g., by directly providing help or public recognition. If the planner targets two individuals who are neighbors and encourages both to work harder, one individual's increase in effort reduces his neighbors' incentives for effort. That is, these two parts of the intervention work against each other, potentially wasting intervention resources. Thus, the planner's intervention may want to move neighbors' incentives in *opposite* directions. In this case, we show that the best policy divides agents into two groups, the positively and negatively targeted; this is done in such a way that there are many links between oppositely targeted agents, and few links between agents whose incentives are altered in the same direction. It turns out that, parallel to the complements case, the welfare-optimal incentive targeting is proportional to an eigenvector of the network—in this case, the one associated to its *least* eigenvalue.

A key idea above is that the network of strategic interactions causes some interventions to be internally self-reinforcing, while others are self-attenuating. Our main technique is a general approach for identifying and separating these, and using that to design interventions. We now sketch the mathematical ideas involved in deriving these results.

Agents' private incentives to put in effort can be represented as a vector of real numbers—marginal benefits that each agent derives from the activity. An intervention is a change to this vector. In setting (1), the planner's problem is to choose the best profile of marginal benefits subject to the constraint of having a total *adjustment cost* not exceeding a resource constraint. In settings (2) and (3), the planner affects the distribution of individuals' idiosyncratic incentives: for example, she may, at a cost, reduce common shocks.[1] In both settings, as we have argued, it will matter which changes to the incentives have positive feedback and are amplified. To find these, we choose particular coordinates for the vector space of possible interventions: an orthonormal basis that is obtained by diagonalizing the matrix of strategic interactions. Each basis vector can be seen as a "basic" intervention. The basis has two key features: (i) the vectors can be ordered in terms of how much strategic feedback or amplification they induce; and (ii) interventions along different basis vectors are "separable" or "independent" in a certain sense, which makes it easy to trace their impact on outcomes.

[1]For some recent studies on how the structure of volatility affects strategic interactions, see, e.g., Angeletos and La'O (2013), Bergemann and Morris (2013), and Bergemann, Heumann, and Morris (2015).

In the case of strategic complements, the basis vector that achieves maximum amplification is an eigenvector corresponding to the largest eigenvalue of the network. In the case of strategic substitutes, the basis vector that achieves maximum amplification is an eigenvector corresponding to the *smallest* eigenvalue of the network. We use this to derive the results sketched above. The eigenvectors that play a key role in these characterizations are related to network measures identified as being important in the existing network literature, in particular Ballester, Calvó-Armengol, and Zenou (2006), Bramoullé, Kranton, and d'Amours (2014) and Acemoglu, Ozdaglar, and Tahbaz-Salehi (2016). We discuss the relation to these papers and others in Section 8.

The decomposition that plays the key role in the arguments has an important interpretation as the *principal component decomposition* of the network adjacency matrix. The basis vector that turns out to be most important in the strategic complements case is the one that, in a suitable sense, best summarizes individuals' levels of connectedness in the whole network. Thus it is called the first principal component. In the case of strategic substitutes, the key vector is one that is, in a sense, the "least representative" of overall social interactions, and the most sensitive to local structure. Thus it is the last principal component. Especially for this latter case, we also give graph-theoretic interpretations, which show that the eigenvector relevant for strategic substitutes encodes information about "how bipartite" the network of interaction is. When the network is actually bipartite, the eigenvector precisely identifies the two sides of the network, i.e., the two disjoint maximal independent sets of the network. To summarize, the principal component interpretation of eigenvectors allows us to relate the results on optimal targeting to network structure.

In general, our approach to finding optimal interventions focuses on the singular value decomposition (SVD) of the adjacency matrix, which is the right generalization, in our problem, of diagonalizing the adjacency matrix. This reason for the SVD's usefulness is that it offers a convenient way to write quadratic forms in the equilibrium actions. Quadratic forms arise naturally in the economic problems we have been discussing. Generally, it will not be only the first or last principal component that matter for optimal targeting. However, the degree to which different components figure in a planner's optimal intervention can generally be nicely ordered, in a way that depends on the strategic structure of the problem at hand and the planner's objective. To illustrate the versatility of our approach, we propose and solve two related intervention problems: one, minimizing volatility of aggregate investment (setting (2) above) and two, maximizing consumer and producer surplus in a supply chain (setting (3) above).

The structure of the optimal interventions can, in general, turn out to be complex: the right targeting scheme may depend not only on network structure, but on other details of the problem, such as the status quo incentives. A natural question is: under what circumstances can we say that the optimal intervention will be *simple*, with the relative degree of focus on different agents determined *only* by the network, and independent of other details? Focusing on the basic network game of setting (1), we show that this occurs for large enough budgets. The network structure determines how large the budget must be for interventions to be simple. Thus we describe which networks admit particularly simple interventions.

Research over the past two decades has deepened our understanding of the empirical structure of networks and how networks affect human behavior. This naturally leads to a study of how policy interventions can effectively exploit network structure, thereby economizing on scarce resources. Our paper contributes to a broad and exciting body of research, spread across economics, sociology, public health, marketing, and computer science, among other fields, which studies network interventions. In economics, this work includes Ballester, Calvó-Armengol, and Zenou (2006), Banerjee, Chandrashekhar, Duflo, and Jackson (2016), Belhaj and Deroian (2017) Bloch and Querou (2013), Candogan, Bimpikis, and Ozdaglar (2012), Demange (2017), Dziubinski and Goyal (2017), Fainmesser and Galeotti (2017), Galeotti and Goyal (2009), and Galeotti and Rogers (2013), as well as many other papers.[2]

The novelty of the paper in the context of this literature lies in three contributions. The first is a unified approach to a class of intervention problems. Our planner's problems are rich in several ways: they allow for a variety of planner's objectives (e.g., total welfare, aggregate volatility, various surpluses); they involve different applications (educational effort, investment, pricing); and they allow both strategic complements and substitutes. Across these problems, the planner's concerns are summarized by closely related quadratic forms that the interventions seek to maximize. The second contribution is in relating the optimal solutions to network structure using the singular value decomposition—and the corresponding analysis of principal components determined by the adjacency matrix. As we have said, all principal components may, in general, play a role in an optimal targeting scheme. The main results show that the degree of emphasis on various components is ordered according to their eigenvalues, even when the intervention is complex. The third contribution is that, in the classical setting of a network game, we characterize when a planner with a welfare objective will have a *simple* targeting policy, focusing on the first or the last principal component. The conditions involved depend on statistics of the network's eigenvalues that capture subtle aspects of its large-scale structure. In this result, and in others, the network statistics and centrality measures that play a key role in our characterizations are, in some cases, standard ones or familiar from recent work; in other cases they are distinct and highlight new aspects of network structure as being economically significant.

The rest of the paper is organized as follows. Section 2 presents the optimal intervention-targeting problem in a canonical network game. Section 3 illustrates, informally, optimal

[2]We refer to Zenou (2016) for a recent survey on interventions in network games. Prominent contributions in related disciplines include Rogers (1983), Feick and Price (1987), Borgatti (2006), Kempe, Kleinberg, and Tardos (2003), and Valente (2012).

interventions in two simple networks. Section 4 sets out notation and basic facts about the singular value decomposition and presents its application to the network game. In Section 5 we provide general results for the optimal targeting intervention problem. Section 6 presents and solves intervention problems for a planner seeking to minimize volatility in an investment game, and to maximize consumer and producer surplus in a supply chain where producers play a pricing game. Section 7 discusses how the main results can be generalized when we relax the assumptions relating to the adjacency matrix, the costs of intervention, and linear-quadratic payoffs. Section 8 relates our work to existing literature on networks games. Section 9 contains concluding remarks. Appendix A contains omitted proofs of some of the propositions.

2. Basic model

There is a set of individuals $\mathcal{N} = \{1, \dots, n\}$ with $n \geq 2$; the individuals are typically indexed by i. Individual i chooses an action $a_i \in \mathbb{R}$, simultaneously with others; the vector of actions across all individuals is denoted $\boldsymbol{a} \in \mathbb{R}^n$. The payoffs to individual i given an action profile $\boldsymbol{a}$ are:

$$W_i(\boldsymbol{a}) = b_i a_i - \frac{1}{2}a_i^2 + \beta a_i \sum_{j \in \mathcal{N}} g_{ij} a_j.$$

In this formulation, $b_i \in \mathbb{R}$ denotes each individual's marginal benefit from his own action; the corresponding vector across all individuals is denoted by $\boldsymbol{b} \in \mathbb{R}^n$. The weighted, directed network with adjacency matrix $\boldsymbol{G}$ has directed links (i, j) with weights g_{ij}; it is a representation of the strategic interactions.

When $\boldsymbol{G}$ is a nonnegative matrix, the parameter β captures the direction of strategic interdependencies. If $\beta > 0$, then actions are strategic complements; if $\beta < 0$, then actions are strategic substitutes.[3] As Ballester et al. (2006) observed, any (pure strategy) Nash equilibrium action profile $\boldsymbol{a}$ satisfies:

$$[\boldsymbol{I} - \beta \boldsymbol{G}]\boldsymbol{a} = \boldsymbol{b}. \tag{1}$$

If the matrix is invertible, the unique Nash equilibrium of the game can be characterized by:

$$\boldsymbol{a} = [\boldsymbol{I} - \beta \boldsymbol{G}]^{-1}\boldsymbol{b}. \tag{2}$$

We will maintain throughout, unless stated otherwise, a standard assumption:

Assumption 1. The spectral radius of $\beta\boldsymbol{G}$ is less than 1.[4]

[3] We do not need the assumption of nonnegative $\boldsymbol{G}$, though the remarks on the interpretation of β show it can be a helpful case to think about.

[4] The spectral radius of a matrix is the maximum of its eigenvalues' absolute values.

This ensures existence of the inverse in (2), and also the uniqueness and stability of the Nash equilibria (Bramoullé et al., 2014).[5]

2.1. **A network intervention problem.** A *status quo* vector of individuals' marginal benefits $\widehat{\boldsymbol{b}}$ is given. The planner wishes to maximize aggregate utility of individuals and can modify, at an *adjustment cost*, every individual's marginal benefit by changing the status quo $\widehat{\boldsymbol{b}}$ to $\boldsymbol{b}$. The timing of the intervention is as follows: the planner moves first and chooses her intervention, and then individuals simultaneously choose actions. Formally, the incentive-targeting problem is:

$$\begin{aligned} &\max_{\boldsymbol{b}} \sum_{i \in \mathcal{N}} W_i(\boldsymbol{a}) \\ &\text{s.t. } \boldsymbol{a} = [\boldsymbol{I} - \beta \boldsymbol{G}]^{-1} \boldsymbol{b}, \\ &\qquad K(\boldsymbol{b}; \widehat{\boldsymbol{b}}) = \sum_{i \in \mathcal{N}} \left(b_i - \widehat{b}_i \right)^2 \leq C, \end{aligned} \tag{IT}$$

where C is the given resource constraint or budget. Thus, the planner is maximizing welfare, subject to agents' playing an equilibrium (given her intervention), and subject to adjustment costs not exceeding the budget.

2.2. **Discussion of assumptions.** The formulation of the adjustment cost $K(\boldsymbol{b}; \widehat{\boldsymbol{b}})$ reflects the idea that the planner faces increasing marginal costs as she seeks to make larger changes in individuals' incentives. We extend the analysis to more general cost functions, and to nonlinear strategic interactions, in Section 7.

Note that the planner may intervene either to encourage or discourage action, i.e., increasing or decreasing b_i relative to the status quo of $\widehat{b}_i$, and that both types of interventions are costly. It is natural to think of the available interventions as changes to the environment that make the action in question more or less appealing. The model is *not* suited to cases where the marginal benefit b_i is, for example, a wage set by our planner, so that the planner is providing all the incentive to take the action in the first place; in that case, the costs would have a different form (we elaborate on possible applications in this direction in Section 9). Instead, it is designed to capture cases where incentives that the planner takes as given set status quo $\widehat{\boldsymbol{b}}$, and the planner can use certain controls to modify them. For instance, suppose a_i is the level of effort toward studying invested by pupils while they are at school; there are also some outside option activities such as, for concreteness, extracurricular clubs. Many of the reasons to pursue each activity are not within a planner's (e.g., school principal's) control, but rather have to do with students' own preferences and future outcomes. Those set the

[5]The game we have presented is an instance of a linear-quadratic game played on a network; papers that study such games include Ballester et al. (2006), Bramoullé et al. (2014), and Goyal and Moraga-Gonzalez (2001). For a survey of research in games on networks, see Bramoulle and Kranton (2016) and Jackson and Zenou (2015).

status quo $\widehat{\boldsymbol{b}}$. Instruments available to the planner for modifying it include facilitating either activity: for example, she can offer some students supplementary tutoring, or offer others improved coaching. Depending on which kind of intervention is done, these interventions tend to increase or decrease b_i, the marginal benefit (relative to an outside option) of studying. Assuming our planner starts with the least costly instruments to achieve her desired influence on a given individual and works up to more costly ones, the marginal costs of changing the player's incentives will be increasing in the magnitude of the change.

For some applications, it may be natural to assume that actions—such as studying effort or prices—can take on a positive value only. This can be reflected in the constraints of the problem. For example, we may study a setting where $\widehat{\boldsymbol{b}}$ is large and C is such that all actions are positive at any solution the planner contemplates.[6]

3. An example

We give a concrete illustration of the main insights of our analysis by presenting the solution to the optimal intervention problem in the case of two networks: a small random network[7] and a circle. These are shown in Figure 1; in each case, the matrix $\boldsymbol{G}$ is the adjacency matrix of the undirected graph, with $g_{ij} = g_{ji} = 1$ if and only if i and j are linked. For the case of strategic complements we set $\beta = 0.1$, and for strategic substitutes we set $\beta = -0.1$.[8] It is easier to explain the optimal intervention policy and the intuitions behind it when the intervention budget is large (in a sense we will make precise later), so we set $C = 1000$.

Figures 1 and 2 present optimal interventions for the different settings. Define $\Delta b_i = b_i^* - \widehat{b}_i$ as the change to marginal benefits made at the optimal intervention. The size of a node corresponds to the magnitude of that change: i.e., $|b_i^* - \widehat{b}_i|$. The color reflects the direction of change: if node i is green (red) it means that the intervention has increased (decreased) the attribute from the initial value of $\widehat{b}_i$. Tables 1-4 present data on the initial b_i, the optimal intervention Δb_i, as well as on the associated change in action Δa_i and change in utility ΔW_i.

Tables 1-4 also present the "first" and "last" eigenvectors of the network $\boldsymbol{G}$. To explain what these are in terms of the network, we introduce some notation. Since each $\boldsymbol{G}$ considered in the example is symmetric, it is diagonalizable, and we can write $\boldsymbol{G} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{\mathsf{T}}$, where (i) $\boldsymbol{\Lambda}$ is a diagonal matrix, whose diagonal has the eigenvalues of $\boldsymbol{G}$ ordered from greatest to least as real numbers, and (ii) $\boldsymbol{U}$ is an orthogonal matrix. The l^{th} column of $\boldsymbol{U}$, which we call $\boldsymbol{u}^l$, is a normalized eigenvector of $\boldsymbol{G}$ associated to the eigenvalue $\Lambda_{ll} = \lambda_l$. In the examples of this section, all the eigenvalues are distinct and the eigenvectors are uniquely determined.

[6]Of course, one could also impose nonnegativity of b_i as a separate constraint, not via C: this would make some arguments more complicated, though the basic forces we will identify would be present in that world as well.

[7]An Erdős-Rényi graph with $p = 0.5$.

[8]All our graphs satisfy the spectral radius condition.

(A) Random network (B) Circle network

FIGURE 1. Optimal intervention with strategic complements

node i	initial benefit $\widehat{b}_i$	first eigenvector u_i^1	scaled intervention $\Delta b_i/\sqrt{C}$	action change Δa_i	welfare change Δw_i
1	0.58	0.51	0.51	23.77	289.92
2	−0.16	0.32	0.31	14.68	110.59
3	0.38	0.40	0.40	18.63	182.49
4	−1.68	0.25	0.24	11.50	69.46
5	0.82	0.24	0.24	11.18	68.90
6	1.81	0.39	0.39	18.09	180.58
7	0.82	0.24	0.24	11.18	68.90
8	−0.19	0.39	0.39	18.08	180.25

TABLE 1. Targeting in a random network, $C = 1000$, $\beta = 0.1$. The eigenvector $\boldsymbol{u}^1$ to which the actual intervention is being compared is the *first* (largest-eigenvalue) eigenvector of $\boldsymbol{G}$.

3.1. **Strategic complements.** Figure 1 illustrates optimal intervention in the case of strategic complements. Optimal intervention entails targeting the nodes in proportion to the corresponding entry in the eigenvector $\boldsymbol{u}^1$. This vector is necessarily positive entrywise.[9] Its entries are individuals' *eigenvector centralities*. The eigenvector centrality of an individual measures his overall or global level of connectedness. It is characterized[10] by the condition that it is entrywise nonnegative and, for all i, we have $u_i^1 = c\sum_j g_{ij}u_j^1$ for some constant $c > 0$. That is, a node is highly central in proportion to the connections (g_{ij} it has), weighted by the centrality of the partners in those connections.

[9]Since both networks have a symmetric and non-negative $\boldsymbol{G}$ and are connected, by the Perron-Frobenius Theorem, $\boldsymbol{u}^1$ is entrywise positive; indeed, this vector is the Perron vector of the matrix.
[10]Given connectedness of $\boldsymbol{G}$, which holds in these examples.

node i	initial benefit $\widehat{b}_i$	first eigenvector u_i^1	scaled intervention $\Delta b_i/\sqrt{C}$	action change Δa_i	welfare change Δw_i
1	0.71	0.35	0.36	14.19	111.26
2	0.00	0.35	0.34	13.73	96.31
3	0.71	0.35	0.36	14.18	110.96
4	0.00	0.35	0.34	13.73	96.24
5	0.71	0.35	0.36	14.19	111.11
6	0.01	0.35	0.34	13.77	96.95
7	0.72	0.35	0.36	14.22	111.73
8	0.00	0.35	0.34	13.76	96.61

TABLE 2. Targeting in a circle network, $C = 1000$, $\beta = 0.1$. The eigenvector $\boldsymbol{u}^1$ to which the actual intervention is being compared is the *first* (largest-eigenvalue) eigenvector of $\boldsymbol{G}$.

Under strategic complementarities, and a large budget, it is by adjusting agents' incentives in proportion to their eigenvector centralities that the planner best amplifies her intervention to maximize aggregate welfare. In the random network, agent 1 has the highest eigenvector centrality and the change Δb_1 chosen by the planner is 16.05, which is very close to $\sqrt{C}\boldsymbol{u}_1^1$; by contrast, node 7 has the lowest eigenvector centrality and Δb_7 is only 7.65, which is very close to $\sqrt{C}\boldsymbol{u}_7^1$. In the circle network nodes have the same structural positions, and thus the same centralities. Thus, any heterogeneity in targeting is only due to differences in $\widehat{b}_i$.[11] As Figure 1b illustrates, these initial differences are less important because of the large budget: the magnitude of the intervention and the consequent change in action is similar across nodes.

3.2. **Strategic substitutes.** Figure 2 illustrates optimal intervention in the case of strategic substitutes. Now the intensity of intervention varies approximately in proportion to the "last" eigenvector, $\boldsymbol{u}^n$. This entails raising the b_i for some nodes and lowering the b_i of others. Figure 2b shows that in the circle network the optimal intervention is to raise b_i's of nodes $\{1,3,5,7\}$ from their initial levels, and to lower those of nodes $\{2,4,6,8\}$. This leads, in turn, to an increase in the actions of nodes $\{1,3,5,7\}$ and a fall in the action of nodes $\{2,4,6,8\}$. Figure 2a shows that a combination of positive and negative interventions is involved in the random network, and here too the interventions track the "last" eigenvector.

To see why this happens, it is instructive to examine the nature of best replies: an increase in b_i raises a_i and this exerts, due to the strategic substitutes property, a downward pressure on neighbor j's action, a_j. A smaller a_j in turn pushes up a_i further, and that lowers a_j even more, and so forth, until we reach a new equilibrium configuration. This process is amplified if we simultaneously increase b_i and decrease b_j. On the other hand, if we were to raise b_i

[11]In both the networks we study, we choose a nonuniform initial vector $\widehat{\boldsymbol{b}}$. This illustrates that the conclusions about targeting being in line with certain eigenvectors are not reliant on any particular structure of the $\widehat{\boldsymbol{b}}$.

and b_j simultaneously, then the pressure toward a higher effort by i and j would tend to cancel each other; that would be wasteful.

The smallest eigenvalue of $\boldsymbol{G}$ and the associated eigenvector $\boldsymbol{u}^n$ are familiar objects in matrix algebra. They contain information about the local structural properties of the network. This information is useful in determining the bipartiteness of a graph and its chromatic number.[12] Consider the circle network in the example; that is a bipartite graph. The associated eigenvector $\boldsymbol{u}^n$ precisely determines the way to partition the nodes of the graph to get the two maximally independent sets. The random network in the example is not bipartite, but one can see that the individuals that have been targeted in the same direction have very few connections among themselves, whereas most of the links are across individuals who have been targeted in opposite directions.

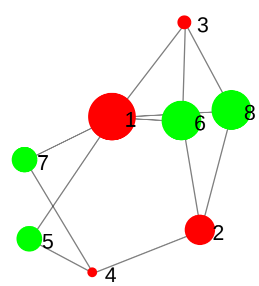

(A) Random Network

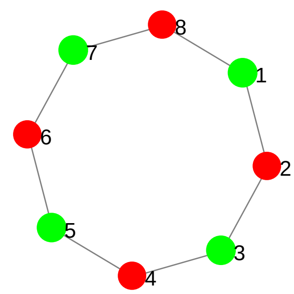

(B) Circle Network

FIGURE 2. Optimal intervention with strategic substitutes

The rest of the paper formalizes and generalizes these insights. Some of the network statistics that turn out to be relevant for solving the intervention problem have been studied in the literature of networks. We discuss these relations in Section 8.

4. Analysis of the game via the singular value decomposition

This section recalls basic concepts that we use to analyze the optimal intervention policy for general network games. Section 4.1 introduces our notation for the *singular value decomposition* (SVD) of an arbitrary matrix $\boldsymbol{M}$ and the associated notion of principal components. When applied to the matrix $\boldsymbol{M}(\boldsymbol{G}) = \boldsymbol{I} - \beta\boldsymbol{G}$, this decomposition helps us identify the highest-feedback interventions. This method—focusing on the SVD of $\boldsymbol{I} - \beta\boldsymbol{G}$—is

[12] Desai and Rao (1994) characterize the smallest eigenvalue of a graph and relate it to the degree of bipartiteness of a graph. Alon and Kahale (1997) demonstrate that the last eigenvector of a graph corresponds to an approximate a coloring of the underlying graph, i.e., a labeling by a minimal set of integers to nodes such that no neighboring nodes share the same label.

node i	initial benefit $\widehat{b}_i$	last eigenvector u_i^n	scaled intervention $\Delta b_i/\sqrt{C}$	action change Δa_i	welfare change Δw_i
1	0.58	−0.55	−0.54	−22.70	265.29
2	−0.16	−0.34	−0.33	−13.98	100.59
3	0.38	−0.14	−0.13	− 5.75	15.84
4	−1.68	−0.09	−0.09	− 3.63	6.35
5	0.82	0.27	0.26	11.47	72.02
6	1.81	0.44	0.44	18.49	187.21
7	0.82	0.27	0.26	11.47	72.02
8	−0.19	0.44	0.44	18.48	186.77

TABLE 3. Targeting in a random network, $C = 1000$, $\beta = -0.1$. The eigenvector $\boldsymbol{u}^n$ to which the actual intervention is being compared is the *last* (smallest-eigenvalue) eigenvector of $\boldsymbol{G}$.

node i	initial benefit $\widehat{b}_i$	last eigenvector u_i^n	scaled intervention $\Delta b_i/\sqrt{C}$	action change Δa_i	welfare change Δw_i
1	0.71	0.35	0.36	14.21	111.57
2	0.00	−0.35	−0.34	−13.76	96.74
3	0.71	0.35	0.36	14.20	111.25
4	0.00	−0.35	−0.34	−13.74	96.43
5	0.71	0.35	0.36	14.17	110.8
6	0.01	−0.35	−0.34	−13.73	96.11
7	0.72	0.35	0.36	14.20	111.44
8	0.00	−0.35	−0.34	−13.77	96.84

TABLE 4. Targeting in a circle network, $C = 1000$, $\beta = -0.1$. The eigenvector $\boldsymbol{u}^n$ to which the actual intervention is being compared is the *last* (smallest-eigenvalue) eigenvector of $\boldsymbol{G}$.

general and works for *any* $\boldsymbol{G}$ and any β. To relate the results to the structure of the network $\boldsymbol{G}$ in a familiar way, we often focus on the special case of a *symmetric* matrix of interaction, such as those seen in the examples of Section 3. In that case, the SVD of $\boldsymbol{M}(\boldsymbol{G})$ boils down to the diagonalization of $\boldsymbol{G}$, and we develop the implications for the game in Sections 4.2 and 4.3. Section 7 is devoted to stating the more general forms of the main results.

4.1. **Singular values and principal components: Notation and definitions.** Consider any $n \times m$ matrix $\boldsymbol{M}$ with real entries. A singular value decomposition (SVD) of $\boldsymbol{M}$ is defined to be a tuple $(\boldsymbol{U}, \boldsymbol{S}, \boldsymbol{V})$ satisfying:

$$\boldsymbol{M} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\mathsf{T}}, \tag{3}$$

where:

(1) $\boldsymbol{U}$ is an orthogonal $n \times n$ matrix whose columns are eigenvectors of $\boldsymbol{M}\boldsymbol{M}^{\mathsf{T}}$;

(2) $\boldsymbol{V}$ is an orthogonal $m \times m$ matrix whose columns are eigenvectors of $\boldsymbol{M}^\mathsf{T}\boldsymbol{M}$;
(3) $\boldsymbol{S}$ is an $n \times m$ matrix with all off-diagonal entries equal to zero and non-negative diagonal entries $S_{ll} = s_l$, which are called *singular values* of $\boldsymbol{M}$.

It is a standard fact that an SVD exists.[13]

We can view $\boldsymbol{M}$ as a map $M : \mathbb{R}^m \to \mathbb{R}^n$ sending column vectors in $\mathbb{R}^m$ to column vectors in $\mathbb{R}^n$. Then the matrices $\boldsymbol{V}$ and $\boldsymbol{U}$ can be seen as bases for the domain and codomain, respectively, under which M is represented by a diagonal matrix. The columns of $\boldsymbol{U}$ are called *left-hand singular vectors* of $\boldsymbol{M}$, and the columns of $\boldsymbol{V}$ are called *right-hand singular vectors*. The *l^{th}-ranked singular value* of a matrix $\boldsymbol{M}$ is defined to be the l^{th}-largest, and the l^{th}-ranked singular vector on the left or right is the corresponding column of $\boldsymbol{U}$ or $\boldsymbol{V}$, respectively.

For any vector $\boldsymbol{x} \in \mathbb{R}^m$, let $\underline{\boldsymbol{x}} = \boldsymbol{V}^\mathsf{T}\boldsymbol{x}$ denote the vector $\boldsymbol{x}$ written in the basis of the SVD, and similarly, for $\boldsymbol{y} \in \mathbb{R}^n$, let $\underline{\boldsymbol{y}} = \boldsymbol{U}^\mathsf{T}\boldsymbol{y}$. The basis of the SVD is one in which the map corresponding to $\boldsymbol{M}$ is particularly nice: it simply dilates some components and contracts others, according to the magnitudes of the singular values:

$$\underline{y}_l = s_l \underline{x}_l \qquad \text{for } l \in \{1, 2, \ldots, n\}.$$

An important application and interpretation of the SVD is *principal component analysis*. We can think of the columns of $\boldsymbol{M}$ as m data points. The first principal component of $\boldsymbol{M}$ is defined as the n-dimensional vector that minimizes the sum of squared distances to the actual $\boldsymbol{M}$. The first principal component can therefore be thought of as a fictitious column that best summarizes the data set $\boldsymbol{M}$. To characterize the other principal components, we orthogonally project all columns of $\boldsymbol{M}$ off this vector and repeat this procedure. A well known result is that the left singular vectors of $\boldsymbol{M}$ are, indeed, the principal components of $\boldsymbol{M}$; a singular value quantifies the variation explained by the respective principal component. When we refer to the l^{th} principal component of $\boldsymbol{M}$ we mean the l^{th}-ranked left singular vector of $\boldsymbol{M}$. From now on, we will refer to $\underline{x}_l$ as the projection of it onto the lth principal component, or the magnitude of it in that component.

4.2. **A special case: Symmetric $\boldsymbol{G}$.** Assume that the matrix $\boldsymbol{G}$ is *symmetric*,[14] i.e. that $\boldsymbol{G}^\mathsf{T} = \boldsymbol{G}$. The usefulness of this assumption is brought out in the following statement (see, e.g., Meyer (2000)).

Fact 1. If $\boldsymbol{G}$ is symmetric and Assumption 1 holds, then there is a SVD of $\boldsymbol{M} = \boldsymbol{I} - \beta\boldsymbol{G}$ with $\boldsymbol{U} = \boldsymbol{V}$. This SVD corresponds to a diagonalization $\boldsymbol{G} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^\mathsf{T}$ of $\boldsymbol{G}$ satisfying the following conditions:

[13] A standard references on the SVD is Horn and Johnson (2012). We describe, after stating Fact 1, some conventions for selecting a unique SVD in our setting.

[14] Symmetry of $\boldsymbol{G}$ entails that the impact of i's action on j's incentives is the same as that of j's action on i's incentives.

1. $\boldsymbol{\Lambda}$ is a diagonal matrix whose diagonal elements are the eigenvalues of $\boldsymbol{G}$ (which are nonnegative real numbers), ordered from greatest to least;
2. the i^{th} column of $\boldsymbol{U}$ is a real eigenvector of $\boldsymbol{G}$ associated to the eigenvalue of $\boldsymbol{G}$ in the position (i, i) of $\boldsymbol{\Lambda}$;
3. in the SVD of $\boldsymbol{M}$ in which $\boldsymbol{S} = \boldsymbol{I} - \beta\boldsymbol{\Lambda}$.

For a generic symmetric $\boldsymbol{G}$, all the diagonal entries of $\boldsymbol{S}$ are positive, a fact we will sometimes use.[15] The l^{th} eigenvector of $\boldsymbol{G}$, which we denote by $\boldsymbol{u}^l(\boldsymbol{G})$, corresponds to the l^{th} principal component of $\boldsymbol{G}$. The decomposition is uniquely determined up to (i) a permutation that reorders the eigenvalues in $\boldsymbol{\Lambda}$ and correspondingly reorders the columns of $\boldsymbol{U}$; (ii) a sign flip of any column of $\boldsymbol{U}$.

The implication of Fact 1 is that, when $\boldsymbol{G}$ is symmetric and Assumption 1 holds, the SVD of $\boldsymbol{M} = \boldsymbol{I} - \beta\boldsymbol{G}$ can be taken to have $\boldsymbol{U} = \boldsymbol{V}$, and the SVD basis is one in which $\boldsymbol{G}$ is diagonal.[16] For concrete examples of the principal components involved, recall our application of the decomposition in Section 3.

4.3. **Analysis of the game using the SVD: Special case of symmetric $\boldsymbol{G}$.** Substituting the expression $\boldsymbol{G} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{\mathsf{T}}$ into the equilibrium equation (1), we obtain:

$$[\boldsymbol{I} - \beta\boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^{\mathsf{T}}]\boldsymbol{a} = \boldsymbol{b}.$$

Multiplying the LHS and the RHS by $\boldsymbol{U}^{\mathsf{T}}$ gives an analogue of (2) characterizing the solution of the game:

$$[\boldsymbol{I} - \beta\boldsymbol{\Lambda}]\underline{\boldsymbol{a}} = \underline{\boldsymbol{b}} \qquad \iff \qquad \underline{\boldsymbol{a}} = \boldsymbol{S}^{-1}\underline{\boldsymbol{b}} \qquad \iff \qquad \underline{\boldsymbol{a}} = [\boldsymbol{I} - \beta\boldsymbol{\Lambda}]^{-1}\underline{\boldsymbol{b}}.$$

This system is diagonal. Hence, for any $l \in \{1, 2, \ldots, n\}$,

$$\underline{a}_l = \frac{1}{1 - \beta\lambda_l}\underline{b}_l. \tag{4}$$

The equilibrium action in the l^{th} principal component of $\boldsymbol{G}$ is a scaling of the magnitude of $\boldsymbol{b}$ in that principal component.[17] Hence, *changes in $\boldsymbol{b}$ in a given principal component are entirely confined to that component in terms of their effect on actions.* In terms of magnitudes, suppose $\boldsymbol{b}$ changes in the principal component of $\boldsymbol{G}$ corresponding to a high value of $(1 - \beta\lambda_l)^{-1}$. In this case, the change in $\boldsymbol{a}$ is large. It follows that, for a nonnegative $\boldsymbol{G}$, when actions are strategic substitutes an ϵ-increase of $\underline{b}_l$ increases the action $\underline{a}_l$ by less

[15] Assumption 1 implies that $\beta\boldsymbol{\Lambda}$ has no entries larger than 1 in absolute value, so the only case to worry about is where some entries are exactly equal to 1, which is not generic.

[16] Furthermore, if $\beta < 0$ (which corresponds to strategic substitutes if $\boldsymbol{G}$ is nonnegative), then the l^{th}-ranked principal component of $\boldsymbol{M} = \boldsymbol{I} - \beta\boldsymbol{G}$ is the l^{th}-ranked principal component of $\boldsymbol{G}$; in the opposite case of $\beta > 0$, the l^{th}-ranked principal component of $\boldsymbol{M} = \boldsymbol{I} - \beta\boldsymbol{G}$ is the $(n - l + 1)^{\text{th}}$-ranked principal component of $\boldsymbol{G}$.

[17] Note that $\underline{b}_l$ is the magnitude of the orthogonal projection of $\boldsymbol{b}$ onto column l of $\boldsymbol{V} = \boldsymbol{U}$.

than ϵ; under strategic complements, however, an ϵ increase of $\underline{b}_l$ increases the action $\underline{a}_l$ by more than ϵ.

Rewriting in the original coordinates:

$$a_i = \sum_l u_i^l \frac{\underline{b}_l}{1 - \beta\lambda_l}.$$

Thus individual i's action is proportional to how much i is represented in various components (u_i^l) as l ranges across all indices; how large the attribute vector is in those components ($\underline{b}_l$); and the magnification from the corresponding factor, $(1 - \beta\lambda_l)^{-1}$.

The same analysis can actually be extended to all normal matrices $\boldsymbol{G}$ (ones that satisfy $\boldsymbol{G}\boldsymbol{G}^{\mathsf{T}} = \boldsymbol{G}\boldsymbol{G}^{\mathsf{T}}$, and are therefore orthogonally diagonalizable) by grouping eigenvalues that are complex conjugates together in the summation. However, the symmetric case illustrates the key ideas, and we leave the general case, which does not require diagonalizability at all, for Section 7.

5. Targeting incentives to maximize welfare

We are now in a position to state our first main result on optimal targeting of interventions to increase utilitarian welfare in the problem (IT) of Section 2.1. Recall that the planner solves the following optimization problem:

$$\begin{aligned} &\max_{\boldsymbol{b}} \sum_{i \in \mathcal{N}} W_i(\mathbf{a}) \\ &\text{s.t. } \boldsymbol{a} = [\boldsymbol{I} - \beta \boldsymbol{G}]^{-1} \boldsymbol{b}, \\ &\qquad K(\boldsymbol{b}; \widehat{\boldsymbol{b}}) = \sum_{i \in \mathcal{N}} \left(b_i - \widehat{b}_i\right)^2 \leq C, \end{aligned} \tag{IT}$$

where the resource constraint is C, and $\widehat{\boldsymbol{b}}$ is a fixed vector of status quo attributes.

5.1. **The structure of optimal interventions in terms of principal components.** Our first result shows that optimal interventions, in a suitable sense, focus on changing $\boldsymbol{b}$ in some principal components more than in others. Recall that for an arbitrary vector $\boldsymbol{b}$, its tuple of SVD coordinates (equivalently, the same vector written the SVD basis) is denoted by $\underline{\boldsymbol{b}}$. Here $\underline{b}_l$ is the projection of $\boldsymbol{b}$ onto the lth principal component of $\boldsymbol{G}$, corresponding to its lth-largest eigenvalue. The statement of our result pertains to *relative* changes in these components. Thus, for an arbitrary vector $\boldsymbol{b}$, we let:

$$\underset{\sim}{b}_l = \frac{\underline{b}_l - \underline{\widehat{b}}_l}{\underline{\widehat{b}}_l}$$

when these are well-defined (i.e. when the denominators $\widehat{\underline{b}}_l$ are nonzero). The quantity $\underset{\sim}{\boldsymbol{b}}$ describes the *relative* increment in a given $\underline{b}_l$ from its status quo level of $\widehat{\underline{b}}_l$; the increase is expressed as a fraction of the initial level.

Proposition 1. Suppose $\boldsymbol{G}$ is symmetric and Assumption 1 holds. Let $\boldsymbol{b}^*$ be a solution to the incentive-targeting problem (IT) with graph $\boldsymbol{G}$. For a generic[18] $\widehat{\boldsymbol{b}}$, we have that $\underset{\sim}{b}_l^* > 0$ for all l, and that:

(1) If $\beta > 0$ then $\underset{\sim}{b}_l^*$ is weakly decreasing in l;
(2) If $\beta < 0$ then $\underset{\sim}{b}_l^*$ is weakly increasing in l.

The proposition says that, in the relative sense described above, the planner focuses her budget C most on changing the contribution of a particular principal component. This is the one corresponding to λ_1, the largest eigenvalue of $\boldsymbol{G}$, or λ_n, the smallest eigenvalue of $\boldsymbol{G}$, depending on whether $\beta > 0$ or $\beta < 0$, respectively.[19] Moreover, the degree of focus on principal components is monotonic in the corresponding eigenvalues. If $\beta > 0$ the degree of focus on a component with a larger eigenvalue is larger. On the other hand, when $\beta < 0$, it is just the opposite. When $\boldsymbol{G}$ is nonnegative, the former case corresponds to strategic complements and the latter to strategic substitutes.

The idea of the proof is as follows: First, we rewrite the problem (IT) in the coordinates of the SVD:

$$\max_{\underline{\boldsymbol{b}}} \sum_{l \in \mathcal{N}} \alpha_l \underline{b}_l^2 \qquad \text{s.t.} \sum_l [\underline{b}_l - \widehat{\underline{b}}_l]^2 \leq C. \tag{IT-SVD}$$

This transformation uses (i) that orthogonal transformation into the SVD coordinates does not change sums of squares of coordinates, so the constraint inequality remains identical in form; (ii) the magnitude of the equilibrium action in the l^{th} principal component of $\boldsymbol{G}$ is simply a scaling of the magnitude of $\underline{b}_l$ (recall (4)) by a coefficient we call α_l, which depends on a corresponding eigenvalue of $\boldsymbol{G}$. Then we make one more transformation, writing the objective and the constraint equivalently in terms of the *relative* changes, $\underset{\sim}{\boldsymbol{b}}$. After a few steps of simplification, the optimization problem becomes:

$$\max_{\underset{\sim}{\boldsymbol{b}}} \sum_{l \in \mathcal{N}} \alpha_l \widehat{\underline{b}}_l^2 [\underset{\sim}{b}_l + 1]^2 \qquad \text{s.t.} \sum_l \widehat{\underline{b}}_l^2 \underset{\sim}{b}_l^2 \leq C. \tag{IT-SVD-REL}$$

[18]Nonzero in each component.
[19]Recall that for the purpose of indexing eigenvalues and the corresponding eigenvectors, the eigenvalues of $\boldsymbol{G}$ are ordered from greatest or least as real numbers.

From this it is straightforward to argue, using basic optimization theory, that at the optimal solution $\underset{\sim}{\boldsymbol{b}}^*$, the entries $\underset{\sim}{b}_l^*$ are increasing in the corresponding α_l; meanwhile, the α_l are shown to be monotone in the eigenvalues (decreasing in λ_l when $\beta > 0$, and increasing when $\beta < 0$). The details are presented in Section A.1.1 of the appendix.

5.2. **Simple interventions under large budgets.** Beyond what we know from Proposition 1, the details of the planner's intervention may be complex: the details of what exactly $\boldsymbol{b}^*$ is—which combination of the principal components—will depend in subtle ways on details of the network $\boldsymbol{G}$, the strategic parameter β, and the budget C. In this section we ask when the optimal policy is simple, in a sense we make precise.

To facilitate this, we define a standard notion of similarity between two vectors.

Definition. The *cosine similarity* of two non-zero vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ is:

$$\rho(\boldsymbol{x}, \boldsymbol{y}) := \frac{\boldsymbol{x} \cdot \boldsymbol{y}}{\|\boldsymbol{x}\| \|\boldsymbol{y}\|}.$$

This is the cosine of the angle between the two vectors in the plane that $\boldsymbol{x}$ and $\boldsymbol{y}$ define; when this is 1, one vector is a positive scaling of the other. When the cosine similarity is 0, the vectors are orthogonal; when the cosine similarity is -1, one vector is a negative scaling of the other.

Our main result in this section shows that, when the resources available for intervention C are appropriately large, the optimal intervention targets individuals according to a single component of the network, in the sense that the optimal intervention has high similarity to that component and low similarity to all others. Before stating the result, we introduce two network-dependent quantities that play a key role in it.[20]

$$\kappa_+ = \left[\left(\frac{1 - \beta\lambda_2}{1 - \beta\lambda_1} \right)^2 - 1 \right]^{-1} \qquad \text{and} \qquad \kappa_- = \left[\left(\frac{1 - \beta\lambda_{n-1}}{1 - \beta\lambda_n} \right)^2 - 1 \right]^{-1}.$$

Proposition 2. Suppose $\boldsymbol{G}$ is symmetric and Assumption 1 holds. Suppose also that λ_1 and λ_n are *strict* bounds on the other eigenvalues.[21] Let $\boldsymbol{b}^*$ be a solution to the incentive-targeting problem (IT) with graph $\boldsymbol{G}$ and a generic $\widehat{\boldsymbol{b}}$. Then for suitable choice of signs[22] for $\boldsymbol{u}^1$ and $\boldsymbol{u}^n$, we have:

1. Suppose $\beta > 0$. For any $\epsilon > 0$, if $\sqrt{C} > L_+ \cdot \kappa_+/\epsilon$, then $\rho(\Delta\boldsymbol{b}^*, \boldsymbol{u}^1) > 1 - \epsilon$, where $L_+ = \|\widehat{\boldsymbol{b}}\|^4 / 2\underline{\widehat{b}}_1^2$.
2. Suppose $\beta < 0$. For any $\epsilon > 0$, if $\sqrt{C} > L_- \cdot \kappa_-/\epsilon$, then $\rho(\Delta\boldsymbol{b}^*, \boldsymbol{u}^n) > 1 - \epsilon$, where $L_- = \|\widehat{\boldsymbol{b}}\|^4 / 2\underline{\widehat{b}}_n^2$.

[20]Both quantities depend on $\boldsymbol{G}$ through the eigenvalues, but we drop this argument when there is no risk of confusion.
[21]This assumption holds for generic symmetric $\boldsymbol{G}$.
[22]Recall that the eigenvectors $\boldsymbol{u}^1$ and $\boldsymbol{u}^n$ are determined up to a sign flip.

This Proposition yields an immediate corollary, which is the main take-away: the optimal intervention is similar to exactly one principal component.

Corollary 1. Take the assumptions of Proposition 2.

(1) If $\beta > 0$, then as $C \to \infty$, the intervention becomes similar to $\boldsymbol{u}^1$ only; i.e., $\rho(\Delta \boldsymbol{b}^*, \boldsymbol{u}^1)$ tends to 1 and $\rho(\Delta \boldsymbol{b}^*, \boldsymbol{u}^l)$ tends to 0 for $l \neq 1$.
(2) Suppose $\beta < 0$. Then as $C \to \infty$, the intervention becomes similar to $\boldsymbol{u}^n$ only; i.e., $\rho(\Delta \boldsymbol{b}^*, \boldsymbol{u}^n)$ tends to 1 and $\rho(\Delta \boldsymbol{b}^*, \boldsymbol{u}^l)$ tends to 0 for $l \neq n$.

Thus, large budgets C imply that *the optimal intervention strategy is simple to describe: the relative emphasis on different agents will stop changing with β and C, and depends only on a single principal component of the network $\boldsymbol{G}$.* If actions are strategic complements, the optimal intervention $\Delta \boldsymbol{b}^*$ alters the individual incentives (very nearly) proportionally to the first principal component of $\boldsymbol{G}$, namely $\boldsymbol{u}^1(\boldsymbol{G})$. On the other hand, if actions are strategic substitutes, the planner changes the individual incentives (very nearly) in proportion to the last principal component, $\boldsymbol{u}^n(\boldsymbol{G})$.[23] Indeed, a heuristic for computing the optimal intervention is simply to set $\Delta \boldsymbol{b} = \sqrt{C}\boldsymbol{u}^1$ (in the case of strategic complements) or $\Delta \boldsymbol{b} = \sqrt{C}\boldsymbol{u}^n$ (in the case of strategic complements); that clearly exhausts the budget (since vectors $\boldsymbol{u}^l$ are normalized) and puts all the weight on the component we have shown is dominant. In the example of Section 3 note that, in each table, the suitably scaled optimal intervention is indeed very close to the relevant eigenvector, which illustrates that this heuristic gets very close to the optimal targeting scheme.

5.3. **Which networks allow simple interventions?** All the different principal components of the network will play a role in the optimal policy when the budget is not large (Proposition 1). Proposition 2 gives a condition on the size of the budget beyond which this complexity goes away. Loosely, this condition is harder to satisfy when the entries of $\widehat{\boldsymbol{b}}$ are larger or, holding the average of the entries fixed, when they are more variable.[24] In this subsection, we analyze the role that the network plays in this condition, via the coefficient κ_+ for the case of $\beta > 0$ and the coefficient κ_- for the case of $\beta < 0$. The key idea that the discussion will highlight is that in some networks, there are two different structures embedded in a network, each one offering a similar potential for amplifying the effects of interventions. In such networks, the condition on the budget for simplicity will be very stringent—i.e., interventions will *not* be simple for reasonable budgets.

[23] When the initial attributes are zero ($\widehat{\boldsymbol{b}} = \boldsymbol{0}$), we can dispense with the approximations. Assuming $\boldsymbol{G}$ is generic in the sense used in Proposition 2, if $\widehat{\boldsymbol{b}} = \boldsymbol{0}$, then all of C is spent either (i) on increasing $\underline{b}_1$ (if $\beta > 0$), or (ii) on increasing $\underline{b}_n$ (if $\beta < 0$). To see this, set $\widehat{\underline{\boldsymbol{b}}} = \boldsymbol{0}$ in the maximization problem (IT-SVD); note that if the allocation is not monotonic then effort can be reallocated profitably among principal components without changing the cost.

[24] Recall that $\|\frac{1}{n}\widehat{\boldsymbol{b}}\|^2$ is equal to $\left(\frac{1}{n}\sum_i \widehat{b}_i\right)^2$ plus the variance of the entries of the vector $\widehat{\boldsymbol{b}}$.

In the strategic complements case, the simplest obstruction to simple interventions occurs when a network consists of two communities, strongly segregated from each other in terms of externalities. What share of her resources the planner would devote to each of these nearly decoupled components depends not only on the network but on which community's initial conditions are more hospitable to interventions—i.e., the details of $\widehat{\boldsymbol{b}}$. In that case, the network need *not* admit simple—i.e. $\widehat{\boldsymbol{b}}$-independent—characterizations of interventions. On the other hand, networks that are more *cohesive*, in the sense that such segregation is ruled out, *will* admit simple interventions. Formally, in the case of strategic complements, κ_+ and therefore the lower bound, is large when $\lambda_1 - \lambda_2$ is small. The difference $\lambda_1 - \lambda_2$ is called spectral gap and it is known th measure the extent of the segregation we have intuitively described: The network spectral gap is small when the nodes can be divided into at least two communities that have few links between them (Hartfiel and Meyer, 1998; Levin et al., 2009).[25]

The analogous phenomenon in the strategic substitutes case is more subtle, and is not as simple as finding nearly disconnected pieces in the network. We therefore discuss this first within an example of a large random graph with group structure. The random graph is formed as follows:

- Let k be a positive integer and let N consist of four disjoint groups $N_{\mathrm{NW}}, N_{\mathrm{NE}}, N_{\mathrm{SW}}, N_{\mathrm{SE}}$, each of size k.
- For every pair of distinct nodes $i, j \in N$ we independently draw a vertical and horizontal *compatibility*, $C_{\mathrm{V}}(i,j), C_{\mathrm{H}}(i,j) \in \{0,1\}$.
  - The probability that i and j are vertically compatible is $h_{\mathrm{V}} \in (0,1)$ if i and j are in groups with the same vertical (N/S) coordinate, and $1 - h_{\mathrm{V}}$ otherwise.
  - The probability that i and j are horizontally compatible is $h_{\mathrm{H}} \in (0,1)$ if i and j are in groups with the same horizontal (E/W) coordinate, and $1 - h_{\mathrm{H}}$ otherwise.
- If i, j are compatible on *both* dimensions then, i.i.d., a link between them is formed in $\boldsymbol{G}$ with a probability $p(k)$, which controls the density of the graph. For example, the probability that $i \in N_{NW}$ and $j \in N_{NE}$ are linked is $h_{\mathrm{V}}(1 - h_{\mathrm{H}})p(k)$.

In Figure 3(a), we illustrate schematically one parameter configuration. We take β so that $\boldsymbol{M}$ has all eigenvalues positive; as we explain in Appendix A.2, one such choice is to set $\beta(k)$ to be a constant less than 1 times the average degree of the graph.

We describe κ_- when $h_V < h_H < 1/2$—i.e., on both dimensions, there is more association across dimensions than within.[26] In this case, we can give a simple formula for κ_- that holds asymptotically.

[25]See also Golub and Jackson (2012) for more on how the difference $\lambda_1 - \lambda_2$ corresponds to network structure, and in particular segregation and homophily, some other applications of that, and further citations to an extensive literature on this issue in applied mathematics.

[26]Other configurations of the parameters can also be analyzed using the techniques that appear in our proofs.

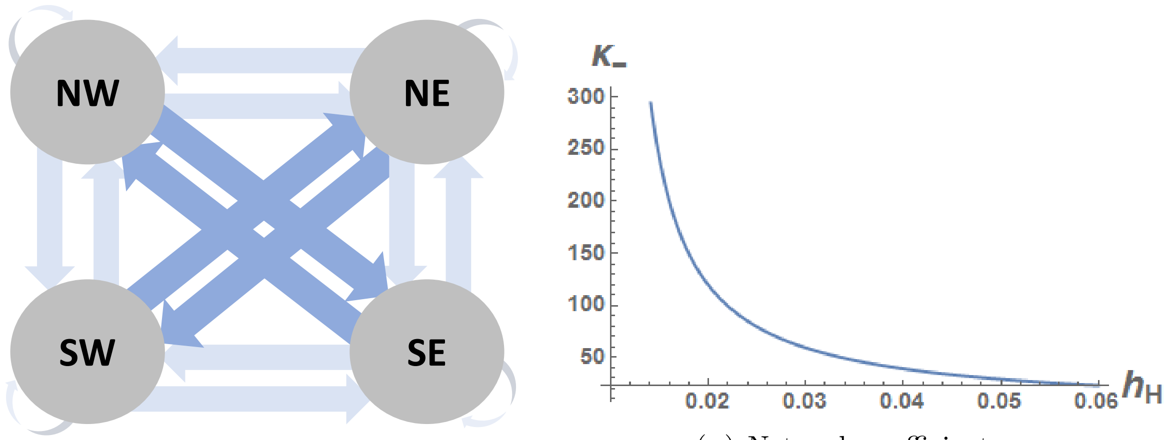

(A) Schematic link structure

(B) Network coefficient κ_-

FIGURE 3. The schematic in (a) depicts the link structure when h_V are very close h_H and near zero, with darkness/thickness indicating the density of the corresponding links. In (b) we fix $\beta = 0.7$ and $h_V = 0.01$ and depict how κ_- varies as h_H approaches 0.01 from above.

Fact 2. If the conditions just stated hold and $kp(k) \gg \log^2 k$ (i.e. the graph is dense enough) then, as $k \to \infty$,

$$\kappa_-(\boldsymbol{G}) \xrightarrow{p} \left[\left(\frac{1 - \beta(2h_H - 1)}{1 - \beta(2h_V - 1)} \right)^2 - 1 \right]^{-1}. \tag{5}$$

This lower bound coefficient grows large when h_H and h_V are very close to each other; see Figure 3(b). In this case, recalling Figure 3(a), there is an East/West and a North/South bipartition of the network. Then either of these bipartitions can serve as the one that best exploits the potential for amplification via strategic substitutes. Therefore, for moderate budgets, the planner's choice will again depend on details of $\widehat{\boldsymbol{b}}$—that will determine which intervention has higher returns. An analogous analysis for the case of strategic complements is relegated to Section A.2 in the appendix.

We now discuss the key idea that generalizes beyond the example. For a fixed value of λ_1, the magnitude of λ_n is, in a way we are about to make precise, a measure of "how bipartite" the graph is: the associated eigenvector $\boldsymbol{u}^n$ is the vector of a given length that is best at assigning different values to neighboring nodes. The magnitude of λ_n is large when it is possible to do this well, and small otherwise.[27] The next-smallest eigenvalue of $\boldsymbol{G}$ captures how good the next-best vector orthogonal to $\boldsymbol{u}^n$ is at doing the same thing. Thus λ_{n-1} and λ_n will be similar if there are two distinct (in the sense of orthogonality) ways to divide the graph, each of which cuts a similar number of edges. (This is exactly what occurs when h_H and h_V are similar, which is what is needed for κ_- to be large in the formula of Fact 2.)

[27]The eigenvalue λ_n is necessarily negative. All the eigenvalues of $\boldsymbol{G}$ are real, and the trace of $\boldsymbol{G}$, which is the sum of the eigenvalues is 0: thus, since λ_1 is positive by the Perron-Frobenius Theorem, λ_n is negative.

To formalize this, we can use a variant of the Courant-Fischer characterization of eigenvalues, which for simplicity we present in the case of a d-regular $\boldsymbol{G}$, one in which the sum of every row and column of $\boldsymbol{G}$ is equal to $d > 0$. This characterization says that:

$$\lambda_n = d - \frac{1}{2} \max_{\boldsymbol{u}\colon \|\boldsymbol{u}\|=1} \sum_{i,j} G_{ij}(u_i - u_j)^2 \qquad \lambda_{n-1} = d - \frac{1}{2} \max_{\substack{\boldsymbol{u}\colon \|\boldsymbol{u}\|=1 \\ \boldsymbol{u}\cdot\boldsymbol{u}^n=0}} \sum_{i,j} G_{ij}(u_i - u_j)^2.$$

Moreover, generically, the eigenvectors $\boldsymbol{u}^n$ and $\boldsymbol{u}^{n-1}$ are the maximizers achieving λ_n and λ_{n-1}, respectively. Thus, $\boldsymbol{u}^n$ achieves the greatest total squared difference of entries u_i across links, and $\boldsymbol{u}^{n-1}$ achieves the greatest among those $\boldsymbol{u}$ orthogonal to $\boldsymbol{u}^n$. If λ_n and λ_{n-1} are close, then these maxima are very similar, meaning there are two splits in the graph that are about equally good from this perspective.

6. Two further applications

This section presents two additional network intervention problems: (i) minimizing aggregate volatility in investment, and (ii) maximizing consumers' and producers' surplus when production occurs in a supply chain.

6.1. **Aggregate volatility.** The network game introduced in Section 2 is related to models of how idiosyncratic shocks contribute to aggregate volatility. A recent strand of research studies how idiosyncratic shocks in production networks affect aggregate volatility of the economy—e.g., Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012). In the production network context, firms are price-taking, $\hat{b}_i$ is a productivity shock to firm i, a_i is the log output of firm i and the production of good i is obtained via combining other goods; g_{ij} indicates how important product j is for the production of good i. See Acemoglu et al. (2016) for a survey of this literature and a formal connection to network games. Other work has studied how private information shapes aggregate volatility in games with a focus, for tractability, on games with linear best replies, e.g., Angeletos and Pavan (2007) and Bergemann, Heumann, and Morris (2015); recent papers (de Martí and Zenou, 2015; Bergemann, Heumann, and Morris, 2017) explore this problem from a network perspective. The underlying game that is studied is one of asymmetric information. Following Golub and Morris (2017), we can extend our complete-information analysis to linear best-response games in which b_i is private information.

Following this latter strand, we interpret actions as levels of investment. The status-quo vector $\widehat{\boldsymbol{b}}$ of exogenous attributes is a shock and its variance-covariance matrix is $\widehat{\boldsymbol{\Omega}}$. The planner, at the time of her intervention, does not observe the realization of $\widehat{\boldsymbol{b}}$, whereas $\widehat{\boldsymbol{b}}$ is common knowledge among the individuals at the time of their investment choice. The planner wishes to minimize the volatility in the level of aggregate investment:

$$\operatorname{Var}\left(\sum_{i\in\mathcal{N}} a_i\right).$$

She can do this by controlling the variances of the shocks to the exogenous attributes. This control comes at a cost: $K(\boldsymbol{\Omega};\widehat{\boldsymbol{\Omega}})$ is the cost of changing the variance-covariance matrix from $\widehat{\boldsymbol{\Omega}}$ to $\boldsymbol{\Omega}$. We make the following assumption on K:

Assumption 2. *K is invariant to rotations of coordinates: for any orthogonal matrix $\boldsymbol{O}$,*

$$K(\widehat{\boldsymbol{\Omega}}+\boldsymbol{D};\widehat{\boldsymbol{\Omega}}) = K(\widehat{\boldsymbol{\Omega}}+\boldsymbol{O}^{\mathsf{T}}\boldsymbol{D}\boldsymbol{O};\widehat{\boldsymbol{\Omega}})$$

and $K(\widehat{\boldsymbol{\Omega}};\widehat{\boldsymbol{\Omega}})=0$.

For example, suppose that the cost $K(\boldsymbol{\Omega};\widehat{\boldsymbol{\Omega}})$ is equal to the reduction in the sum of attribute variances. Then $K(\boldsymbol{\Omega};\widehat{\boldsymbol{\Omega}}) = \left[\sum_i(\omega_{ii}-\widehat{\omega}_{ii})\right]^2 = \left(\operatorname{trace}(\boldsymbol{\Omega}-\widehat{\boldsymbol{\Omega}})\right)^2$. This specification satisfies Assumption 2.[28]

Under this assumption, we will study the variance-minimization problem described above. Formally:

$$\begin{aligned} &\min_{\boldsymbol{\Omega}} \operatorname{Var}\left(\sum_i a_i\right) \\ &\text{s.t. } K(\boldsymbol{\Omega};\widehat{\boldsymbol{\Omega}}) \le C. \\ &\quad \boldsymbol{a} = [\boldsymbol{I}-\beta\boldsymbol{G}]^{-1}\boldsymbol{b}. \end{aligned} \tag{VM}$$

Proposition 3. *Assume $\boldsymbol{G}$ is symmetric and Assumptions 1 and 2 hold. Suppose $\boldsymbol{\Omega}^*$ solves (VM). Consider the variance reduction chosen by the planner in the l^{th} principal component of $\boldsymbol{G}$:*

$$\underline{\Delta}_l = \operatorname{Var}_{\boldsymbol{\Omega}^*}(\boldsymbol{u}^l(\boldsymbol{G})\cdot\boldsymbol{b}) - \operatorname{Var}_{\widehat{\boldsymbol{\Omega}}}(\boldsymbol{u}^l(\boldsymbol{G})\cdot\boldsymbol{b}).$$

Then:

1. *If $\beta>0$, then $\underline{\Delta}_l$ is weakly decreasing in l.*
2. *If $\beta<0$, then $\underline{\Delta}_l$ is weakly increasing in l.*

The proposition says that the amount of variance reduction in a given principal component of $\boldsymbol{G}$ is monotone in the ranking of that principal component. In particular, the first principal component will receive the most focus when investments are strategic complements, and

[28]To see this, note that the trace is the sum of the eigenvalues, and this does not change under conjugacy transformations of the argument. The function on the outside is inessential; we could set $K(\boldsymbol{\Omega};\widehat{\boldsymbol{\Omega}}) = \phi\left(\operatorname{trace}(\boldsymbol{\Omega}-\widehat{\boldsymbol{\Omega}})\right)$, where $\phi:\mathbb{R}\to\mathbb{R}$ is any function such that $\phi(0)=0$. However, in those cases where changing variance more is more costly, then it will be natural to make this function increasing in the absolute value of its argument.

the last principal component will receive the most focus when investments are strategic substitutes.

6.1.1. *Intuition.* We now present the intuition behind the result. The variation in basic incentives comes from shocks to each firm's own idiosyncratic costs and benefits. The planner intervention can can adjust (ex ante) the variances and covariances of these incentives. What should the planner do to reduce the volatility of aggregate investments? Suppose first that decisions among the firms are strategic complements. A perfectly correlated shock that raises all costs is amplified: there are knock-on effects because each firm's decrease in activity reduces the incentives of others to invest, adding to the direct effects of those firms' own shocks. In fact, the type of shock that is most amplifying (at a given size) is the one that is perfectly correlated across agents, with the magnitude of a given agent's shock proportional to the first principal component. This is the dimension of volatility that the planner most wants to reduce. In the case of the circle, this can be visualized with the help of Figure 1(b): in a given shock realization, all nodes are shocked in the same (here, positive) direction.

If decisions are strategic substitutes, then a global shock is not so bad. The first-order response of all firms to an increase in costs, is to decrease investment but that, in turn, makes all firms want to increase their investment somewhat due to the strategic substitutability with their neighbors. Hence, the effect of a global shock is attenuated. The *most* amplifying shock profiles, in this case, are the ones in which neighbors have *negatively* correlated shocks. To stabilize aggregate outcomes, a planner will prioritize combating that particular type of volatility. It turns out that the most concerning type of volatility will be the one that is strongly correlated with the last eigenvector of the system, and this is what the planner will focus most on. In the case of the circle, this can be visualized with the help of Figure 2(b): in a given shock realization, nodes are targeted in an alternating way.

6.1.2. *Illustration in the case of the circle.* Proposition 3 says that the planner will also invest in reducing shocks that are neither the totally correlated volatility represented by Figure 1(b) nor the "most local" volatility represented by Figure 2(b). That is, $\underline{\Delta}_l$ may be nonzero for various $l \neq 1, n$. How do the non-extreme components look, and what is the meaning of the statement that the planner's focus is monotonic in l, the index of the eigenvalue? Some intuition is suggested by Figure 4, which shows the $\boldsymbol{u}^l$ corresponding to $l \notin \{1, n\}$, for the circle.[29] The $l = 2$ component (top left panel of Figure 4), represents a type of volatility that is not perfectly correlated (in that opposite sides of the circle are anti-correlated) but quite correlated locally. This component will receive more focus than any but the $l = 1$ component in the case of strategic complements. Indeed, in the case of strategic complements, the planner's efforts will intervene with decreasing intensity as we progress

[29] This basis is not uniquely determined because the circle has a great deal of symmetry; a generic perturbation would uniquely pin down the eigenvectors, so we do not dwell on the details of the multiplicity here.

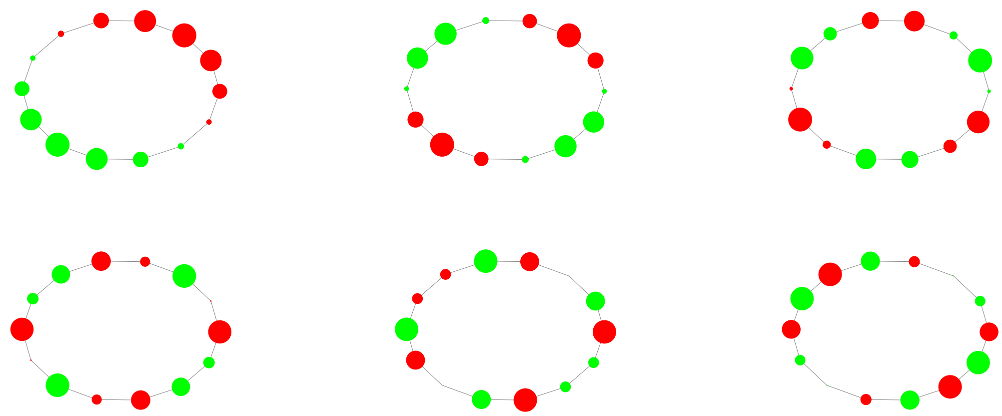

FIGURE 4. For a circle network with 14 nodes, the eigenvectors $\boldsymbol{u}^l$ for $l = 2, 4, 6$ (top row), $l = 8, 10, 12$ (bottom row).

along the sequence depicted in Figure 4. Turning now to the end of the sequence, the $l = 10$ or 12 component (bottom right panel of Figure 4) depicts volatility that is locally quite anti-correlated—the shocks of neighbors are usually opposite—but not quite as strongly as in the last component. These will receive the most focus in the case of strategic substitutes, and generally the planner's interventions will *increase* in intensity as we go along the sequence depicted in the figure.

Another way to summarize what the figure depicts is that the principal components corresponding to smaller (more negative) eigenvalues divide the nodes into more regions where the shock hits with opposite signs. Thus, in the case of strategic complements, a given principal component of volatility is of greater concern for the planner if it has a "longer spatial wavelength." In the case of strategic substitutes, the volatility is of greater concern if it has a "shorter spatial wavelength." These patterns in the eigenvectors, which are illustrated in this simple example, hold more broadly, as studied in Urschel (2018) and references cited there.

6.1.3. *Proof.* The idea of the proof is to consider any solution and to show that if it did not satisfy the conclusion, then it would be possible to find a different variance reduction that does better. The strategy for finding the rearrangement is to study the problem in the eigenvector basis. In this basis, the aggregate volatility is a weighted sum of the variance of each principal component of $\boldsymbol{G}$, and the contribution to aggregate volatility of the variance of each component is monotone in the corresponding eigenvalue. This monotonicity dictates how to rearrange variance reductions to achieve a bigger effect. In particular, we permute them among the eigenvectors. Assumption 2 on the cost function K ensures that the rearrangement we need is feasible without changing the cost. A full proof appears in Section A.3.

6.2. **Pricing in a supply chain.** We consider a pricing game between suppliers embedded in a supply chain. In this application, the intervention alters the variability of marginal costs across suppliers in order to maximize consumer surplus, producer surplus and welfare.

Price formation in networked markets is an active area of research. This research has focused on buyer-seller networks and on networks of intermediaries. To the best of our knowledge, existing work does not address the study of optimal intervention in these markets; for surveys of this literature, see Condorelli and Galeotti (2016), Goyal (2017), and Manea (2016).

We consider a set of *final goods* $\mathcal{F} = \{1, 2, \ldots, F\}$. Final goods are made using the set of *inputs* $\mathcal{N} = \{1, 2, \ldots, N\}$; supplier i produces input $i \in \mathcal{N}$. Following Vives (2001) and Singh and Vives (1984), a representative consumer with quadratic utilities chooses how much to consume of each final good. Given price vector $\boldsymbol{P} = \{P_1, \ldots, P_F\}$, the utility of the consumer is:

$$U(\boldsymbol{Q}) = \sum_{f\in\mathcal{F}} \left(\gamma Q_f - \frac{1}{2}Q_f^2 - P_f Q_f\right).$$

Here, for simplicity, we assume that final goods are independent; the analysis can easily be generalized to the case where final goods can be substitutes and complements in consumption. The consumer's optimization leads to a linear demand for final goods: $Q_f = \gamma - P_f$. The utility of the representative consumer is:

$$U^* = \frac{1}{2}\sum_{f\in\mathcal{F}} Q_f^2.$$

We now describe how inputs are transformed into final goods. Let $\boldsymbol{T}$ be a N-by-F matrix with typical element t_{if}. In order to produce one unit of final good f, firm f requires t_{if} units of each input $i \in \mathcal{N}$; without loss of generality, we set, for each $i \in \mathcal{N}$, $\boldsymbol{t}_i \cdot \boldsymbol{t}_i = 1$.[30] We assume that the final goods markets are competitive and so the price of final good f equals the marginal cost of production of good f.[31] We can thus write $P_f(\boldsymbol{p}) := \sum_{i\in\mathcal{N}} t_{if} p_i$, or, in matrix notation, $P(\boldsymbol{p}) = \boldsymbol{T}^\mathsf{T}\boldsymbol{p}$. The demand of supplier i, which depends on all prices, is:

$$q_i(\boldsymbol{p}) \;=\; \sum_{f\in\mathcal{F}} t_{if} Q_f = \sum_{f\in\mathcal{F}} t_{if}[\gamma - P_f] = \gamma \sum_{f\in\mathcal{F}} t_{if} - \sum_{j=1}^{n} (\boldsymbol{t}_i \cdot \boldsymbol{t}_j) p_j.$$

Thus, the vector of demand for inputs is $q(\boldsymbol{p}) = \boldsymbol{T}Q(\boldsymbol{T}^\mathsf{T}\boldsymbol{p})$. For a given price profile, $\boldsymbol{p}$, the profit of supplier i with a constant marginal cost c_i is:

$$\Pi_i(\boldsymbol{p}) = q_i(\boldsymbol{p})[p_i - c_i].$$

[30] This is a normalization: we choose the relevant units of each input i such that the Euclidean length of each vector $\boldsymbol{t}_i$ is equal to 1.

[31] Constant markups can be added without significantly changing our analysis.

Consider the simultaneous-move pricing game among suppliers, each having profit function Π_i and each taking action p_i. The Nash equilibrium pricing profile $\boldsymbol{p}$ solves the system:

$$[\boldsymbol{I} + \boldsymbol{T}\boldsymbol{T}^{\mathsf{T}}]\boldsymbol{p} = \boldsymbol{b}, \tag{6}$$

where $\boldsymbol{b} = \boldsymbol{c} + \gamma \boldsymbol{T}\boldsymbol{1}$. This is equivalent to system (1) with $\boldsymbol{G} = \boldsymbol{T}\boldsymbol{T}^{\mathsf{T}}$, $\beta = -1$ and the endogenous variables being suppliers' prices, $\boldsymbol{a} = \boldsymbol{p}$. In other words, the pricing game is a special case of the network games we have studied above.

Two observations follow. First, the matrix of interaction across suppliers, $\boldsymbol{G} = \boldsymbol{T}\boldsymbol{T}^{\mathsf{T}}$, is symmetric; furthermore, $\beta = -1$, and so the equilibrium is uniquely defined and stable.[32] Second, the SVD of $\boldsymbol{G} = \boldsymbol{T}\boldsymbol{T}^{\mathsf{T}}$ is related to the SVD of $\boldsymbol{T}$. The SVD of $\boldsymbol{T}$ reads $\boldsymbol{T} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\mathsf{T}}$, where the columns of $\boldsymbol{V}$ and $\boldsymbol{U}$ are the right and left singular vectors of $\boldsymbol{T}$, respectively. It follows that the SVD of $\boldsymbol{G} = \boldsymbol{T}\boldsymbol{T}^{\mathsf{T}}$ is given by $(\boldsymbol{U}, \boldsymbol{\Lambda}, \boldsymbol{U}^{\mathsf{T}})$, where $\boldsymbol{\Lambda} = \boldsymbol{S}\boldsymbol{S}^{\mathsf{T}} = \boldsymbol{S}^2$. Hence, the singular values of $\boldsymbol{G}$ are the square of the singular values of $\boldsymbol{T}$; and the principal components of $\boldsymbol{G}$ turn out to be the left singular vectors of $\boldsymbol{T}$. These can be interpreted as bundles of inputs. In the sense of principal components, they "summarize" the technology $\boldsymbol{T}$ of production.

In this application, we refer to the columns of $\boldsymbol{U}$ as fundamental bundles of inputs and we define $\underline{\boldsymbol{p}} = \boldsymbol{U}^{\mathsf{T}}\boldsymbol{p}$ and $\underline{\boldsymbol{b}} = \boldsymbol{U}^{\mathsf{T}}\boldsymbol{b}$. We can then rewrite the equilibrium price system (6) as follows:

$$\underline{\boldsymbol{p}} = [\boldsymbol{I} + \boldsymbol{\Lambda}]^{-1}\underline{\boldsymbol{b}} \Longleftrightarrow \underline{p}_l = \frac{\underline{b}_l}{1 + \lambda_l}. \tag{7}$$

Now, suppose that the vector $\mathbf{c}$ of marginal costs is common knowledge among market participants, but it is random from the perspective of a planner at the time of the intervention. The variance-covariance matrix of marginal costs prior the intervention is $\widehat{\boldsymbol{\Omega}}$ and the planner can change it to $\boldsymbol{\Omega}$ at a cost $K(\boldsymbol{\Omega} - \widehat{\boldsymbol{\Omega}})$, which satisfies Assumption 2.

We study the optimal choice of $\boldsymbol{\Omega}$ under the constraint $K(\boldsymbol{\Omega} - \widehat{\boldsymbol{\Omega}}) \leq C$, for three objectives: expected consumer surplus $CS(\boldsymbol{p})$, expected producer surplus $PS(\boldsymbol{p})$, and social welfare $SW(\boldsymbol{p}) = CS(\boldsymbol{p}) + PS(\boldsymbol{p})$. It is worth noting right away that a planner who wishes to maximize say, $CS(\boldsymbol{p})$, has preferences for variance in $\boldsymbol{p}$. This is because variance in $\boldsymbol{p}$ generates variance in the quantity consumed of final goods, and the utility of the representative consumer is $\sum_f Q_f^2$. With other specifications of preferences, the consumer's preferences might imply, instead, that it is better to have more stable input prices.

In any case, the change of the three surplus quantities, when we move from $\widehat{\boldsymbol{\Omega}}$ to $\boldsymbol{\Omega}$, turns out to be a weighted sum of the respective changes in the variances of the marginal costs of the fundamental bundles of final goods. The weight associated to the l^{th}-ranked fundamental

[32] Indeed, all eigenvalues of $\boldsymbol{M} = \boldsymbol{I} - \beta\boldsymbol{G}$ are positive: they are one plus the eigenvalues of $\boldsymbol{G}$, which are at most 1 in absolute value by our normalization of $\boldsymbol{T}$. It follows that $\boldsymbol{I} - \beta\boldsymbol{G}$ is invertible, so equilibria are uniquely defined and also stable.

bundle is a function of the l^{th}-ranked singular value. Formally:

$$\mathbb{E}_{\boldsymbol{\Omega}}[CS(\boldsymbol{p})] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}[CS(\boldsymbol{p})] = \frac{1}{2}\sum_l \frac{\lambda_l}{(1+\lambda_l)^2}\left[\underline{\omega}_{ll} - \underline{\widehat{\omega}}_{ll}\right] \tag{8}$$

$$\mathbb{E}_{\boldsymbol{\Omega}}[PS(\boldsymbol{p})] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}[PS(\boldsymbol{p})] = \sum_l \frac{\lambda_l^2}{(1+\lambda_l)^2}\left[\underline{\omega}_{ll} - \underline{\widehat{\omega}}_{ll}\right] \tag{9}$$

$$\mathbb{E}_{\boldsymbol{\Omega}}[SW(\boldsymbol{p})] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}[SW(\boldsymbol{p})] = \sum_l \frac{\lambda_l(1+2\lambda_l)}{2(1+\lambda_l)^2}\left[\underline{\omega}_{ll} - \underline{\widehat{\omega}}_{ll}\right]. \tag{10}$$

where we recall from Section 6.1 that $\underline{\omega}_{ll} = \text{Var}(\boldsymbol{u}^l \cdot \boldsymbol{c})$ under $\boldsymbol{\Omega}$ and $\underline{\widehat{\omega}}_{ll} = \text{Var}(\boldsymbol{u}^l \cdot \boldsymbol{c})$ under $\widehat{\boldsymbol{\Omega}}$.

Proposition 4. Assume K satisfies Assumption 2.

1. Suppose $\boldsymbol{\Omega}^*$ maximizes expected producer surplus or expected total welfare. Then the variance change in the marginal cost of the l^{th} fundamental bundle of final goods is decreasing in l.
2. Suppose $\boldsymbol{\Omega}^*$ maximizes expected consumer surplus. Let the $\bar{l}^{\text{th}}$ fundamental bundle of final goods be such that $\lambda(\bar{l}) \geq 1$ and $\lambda(\bar{l}+1) < 1$.[33] The variance change in the marginal cost of the l^{th} fundamental bundle is increasing in l if $l \leq \bar{l}$ and decreasing otherwise.

Proof of Proposition 4. Consider expression (9) and note that the weight on the variance change of the marginal cost of the l^{th} fundamental bundle is increasing in λ_l. The proof then follows by replicating the proof of Proposition 3. The same arguments apply to the expected welfare. Next, consider the expression (8) for the expected consumer surplus. Note that the weight on the variance change of the marginal cost of the l^{th} fundamental bundle is increasing in λ_l for $\lambda_l \in (0,1)$ and it is decreasing in λ_l for $\lambda_l > 1$. The proof then follows by using the same techniques introduced in the proof of Proposition 3. The three expressions (8)-(10) are derived in Lemma 1 in the Appendix. □

There are two main effects driving the results of Proposition 4. The first is a *pass-through effect* across suppliers. The pricing game is a game of strategic substitutes. Therefore, shocks in marginal costs that alter the price of some suppliers are attenuated by the strategic response of other suppliers, changing how much of the price change a consumer ultimately experiences. This effect is summarized in expression (7) that indicates that shocks are attenuated more along the higher-eigenvalue principal components. The second effect is a *quantity effect.* Any shock to marginal costs is passed through to suppliers' prices and affects the price of final goods, and so the final consumption of the representative consumers. In particular, the equilibrium prices of final goods are:

$$\boldsymbol{P} = \boldsymbol{T}^{\mathsf{T}}\boldsymbol{p} \Longleftrightarrow \underline{\boldsymbol{P}} = \boldsymbol{S}^{\mathsf{T}}\underline{\boldsymbol{p}},$$

[33]Set $\bar{l} = n+1$ if $\lambda(n) > 1$ and set $\bar{l}^{\text{th}} = 0$ if $\lambda(1) < 1$.

where the equivalence follows using the SVD of $\boldsymbol{T}$. Hence, the quantity effects are amplified along the higher-eigenvalue principal components.[34]

When the objective of the planner is to maximize consumer surplus, these two effects are countervailing: the pass-through effect pushes the planner to increase variance in less representative fundamental bundles, whereas the quantity effect pushes her in the direction of the more representative fundamental bundles. In this case, the solution is generally to target variance change on intermediate principal components. When the objective is the producer surplus, the planner only cares about the pass-through effect and so the focus of variance increase is on the main principal component.[35]

7. Generalizations: General $\boldsymbol{G}$, different cost functions, and nonlinear systems

The restrictions to a symmetric $\boldsymbol{G}$ and a linear system of best replies facilitate the simplest presentation of the key ideas in our approach. However, these restrictions will often fail to be satisfied in practical applications. To conclude the formal analysis, we show that many of the insights of the previous sections are valid even when those restrictions are relaxed. Our perspective is that the special cases above are most useful for intuition, while the generalizations here are more robust and portable.

We begin by relaxing the assumption of symmetric $\boldsymbol{G}$; we then study a nonlinear system.

7.1. **Incentive targeting: Beyond symmetric $\boldsymbol{G}$.** In this subsection we extend the analysis of Section 5 to non-symmetric $\boldsymbol{G}$. We maintain the linearity of the best responses, but also relax the functional form assumptions on the objective function.

Individuals take endogenous actions $\boldsymbol{a} = (a_1, \ldots, a_n)$ which, in equilibrium, satisfy:

$$\boldsymbol{a} = \boldsymbol{M}^{-1}\widehat{\boldsymbol{b}},$$

where $\widehat{\boldsymbol{b}}$ is a vector of idiosyncratic characteristics. In the setting of Section 5, let us assume that $\boldsymbol{M} = \boldsymbol{I} - \beta\boldsymbol{G}$ for an arbitrary matrix $\boldsymbol{G}$ such that $\boldsymbol{M}$ is invertible.[36] The planner can intervene and change $\widehat{\boldsymbol{b}}$ to $\boldsymbol{b}$ at a cost, with the aim of maximizing the objective:

$$W(\boldsymbol{a}) = F(\boldsymbol{a}^{\mathsf{T}}\boldsymbol{a}),$$

where F is an increasing function (at least on the domain of $\boldsymbol{a}$ achievable through intervention). When F is a linear function the objective corresponds to the one studied in Section 5.

We use the SVD of $\boldsymbol{M}$ to rewrite the argument of the objective function as a convex combination of the contributions of each of the principal components of $\boldsymbol{M}$ to the argument

[34] If the prices of the l^{th} and $(l+1)^{\text{st}}$ fundamental bundles both increase by ϵ, then the change in $\underline{P}_l$ equals $s_l\epsilon$—larger than the change in $\underline{P}_{l+1}$, which is equal to $s_{l+1}\epsilon$.

[35] This is because the profits of a firm are quadratic in $(p_i - c_i)$.

[36] This amounts to the (generically satisfied) requirement that $1/\beta$ is not an eigenvalue of $\boldsymbol{G}$.

of F, and thus the welfare W. This gives a simple way to describe the marginal benefit of targeting a specific principal component l: that is, the effect of targeting the individuals in the economy proportionally to their representation in the l^{th} principal component. Formally, the SVD of $\boldsymbol{M}$ corresponds to the formula $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\mathsf{T}}$ and therefore, using the transformed coordinates $\underline{\boldsymbol{b}} = \boldsymbol{U}^{\mathsf{T}}\boldsymbol{b}$, we obtain:

$$\boldsymbol{a}^{\mathsf{T}}\boldsymbol{a} = \underline{\boldsymbol{b}}^{\mathsf{T}}[\boldsymbol{S}^{\mathsf{T}}\boldsymbol{S}]^{-1}\underline{\boldsymbol{b}} = \sum_{l=1}^{n} \frac{1}{s_l^2}\underline{b}_l^2.$$

Hence, the objective function is:

$$W = F\left(\sum_{l=1}^{n} \frac{1}{s_l^2}\underline{b}_l^2\right).$$

It is now apparent that the analysis of optimal intervention can be carried out in the same way as in Section 5. The result is that, under the optimal intervention, the degree of focus of the planner to principal component l of $\boldsymbol{M}$ is decreasing in its corresponding singular value s_l.[37]

7.2. **More general costs of intervention in incentive-targeting.** Again we work in the setting of Section 5, now generalizing the cost function. Consider any function $(\boldsymbol{b}; \widehat{\boldsymbol{b}}) \mapsto K(\boldsymbol{b}; \widehat{\boldsymbol{b}})$ and assume the following about it.

Assumption 3. If $\boldsymbol{O}$ is an orthogonal matrix and $\boldsymbol{x} \in \mathbb{R}^n$ then:

$$K(\widehat{\boldsymbol{b}} + \boldsymbol{x}; \widehat{\boldsymbol{b}}) = K(\widehat{\boldsymbol{b}} + \boldsymbol{O}\boldsymbol{x}; \widehat{\boldsymbol{b}}).$$

That is, the costs of interventions are rotationally symmetric, in the sense that they stay fixed when we rotate a given change vector $\boldsymbol{x}$ around $\widehat{\boldsymbol{b}}$.[38] Moreover, assume that costs are increasing as we scale the distance from $\widehat{\boldsymbol{b}}$. Formally,

Assumption 4. If $\boldsymbol{x} \in \mathbb{R}^n$ and $s > 1$ then:

$$K(\widehat{\boldsymbol{b}} + s\boldsymbol{x}; \widehat{\boldsymbol{b}}) > K(\widehat{\boldsymbol{b}} + \boldsymbol{x}; \widehat{\boldsymbol{b}}).$$

Now we can study the following generalization of the incentive-targeting problem:

$$\max_{\boldsymbol{b}} F(\boldsymbol{a}^{\mathsf{T}}\boldsymbol{a}) \quad \text{s.t. } K(\boldsymbol{b}; \widehat{\boldsymbol{b}}) \leq C. \tag{IT-G}$$

[37]The differences from the case in which $\boldsymbol{M} = \boldsymbol{I} - \beta\boldsymbol{G}$ and $\boldsymbol{G}$ is symmetric are as follows: the singular values s_l of $\boldsymbol{M}$ will not be equal $1 - \beta\lambda_l$, where λ_l are eigenvalues of $\boldsymbol{G}$; the singular vectors of $\boldsymbol{M}$ will not be eigenvectors of $\boldsymbol{G}$; and the left and right singular vectors need not be the same.

[38]The quadratic cost function we consider in the basic setting satisfies Assumption 3.

As in Section 5, define:

$$\underset{\sim}{b}_l = \frac{b_l - \widehat{b}_l}{\widehat{b}_l}.$$

Our main result on this is:

Proposition 5. Assume $\boldsymbol{M}$ is invertible and Assumptions 3 and 4 hold. Suppose $\boldsymbol{b}^*$ solves (IT-G). Then for generic $\widehat{\boldsymbol{b}}$, we have that $\underset{\sim}{b}_l^*$ is decreasing in s_l.

This generalizes Proposition 1. There are two key implications. First, as shown in the proof, what is necessary for our conclusion is that the cost of intervention be an increasing function of the Euclidean distance of the new $\boldsymbol{b}$ from the base point $\widehat{\boldsymbol{b}}$; Assumptions 3-4 guarantee this property. Second, the proof of Proposition 1 can easily be modified to dispense with symmetry of $\boldsymbol{G}$. The only cost is that the result must be stated in terms of singular values of $\boldsymbol{M}$ rather than eigenvalues of $\boldsymbol{G}$.

In both the present case, and the previous subsection, the disadvantage of the result being stated in terms of $\boldsymbol{M}$ rather than $\boldsymbol{G}$ is that the singular values of $\boldsymbol{M}$ depend on β. As a result, we can no longer make clean statements about how strategic complements and substitutes differ. Nevertheless, there is still a clean summary of the priorities of the planner in terms of classical invariants of a matrix capturing the strategic interactions. The question of how the singular vectors of $\boldsymbol{M}$ depend on β in the general case is an interesting one which we leave for future work.

7.3. **Nonlinear systems.** In this section we will study the volatility-control problem of Section 6.1 for a nonlinear system. Consider an economic system in which the endogenous investment profile $\boldsymbol{a}$ solves the following system:

$$a_i = f_i(\boldsymbol{h}_i \cdot \boldsymbol{a} + b_i) \text{ for each } i = 1 \ldots n,$$

where f_i is a differentiable function. Fix a solution $\widehat{\boldsymbol{a}}$ of this system corresponding to $\widehat{\boldsymbol{b}}$, and assume that the conditions of the implicit function theorem hold, so that locally the solution is unique. The system is susceptible to mean-zero shocks to the productivity vector $\widehat{\boldsymbol{b}}$, so that $\widehat{\boldsymbol{b}}$ changes to $\widehat{\boldsymbol{b}} + \check{\boldsymbol{x}}$. These shocks are small in a sense we will make precise. The interest is to control aggregate volatility, i.e., the variance of aggregate actions.

Rewrite the system as:

$$\boldsymbol{a} = \boldsymbol{f}(\boldsymbol{H}\boldsymbol{a} + \boldsymbol{b}),$$

where $\boldsymbol{H}$ is the matrix with rows $\boldsymbol{h}_i$. Let $\check{\boldsymbol{a}}$ denote the change to $\boldsymbol{a}$; furthermore we denote by F_i the change of f_i. We see that:

$$\check{\boldsymbol{a}} = \boldsymbol{F}\boldsymbol{H}\check{\boldsymbol{a}} + \boldsymbol{F}\check{\boldsymbol{x}} + O(\|\check{\boldsymbol{x}}\|^2).$$

Define $\boldsymbol{G} = \boldsymbol{F}\boldsymbol{H}$. Then:

$$[\boldsymbol{I} - \boldsymbol{G}]\, \check{\boldsymbol{a}} = \boldsymbol{F}\check{\boldsymbol{x}} + O(\|\check{\boldsymbol{x}}\|^2).$$

We decouple the system by using the SVD of $\boldsymbol{M} = \boldsymbol{I} - \boldsymbol{G}$, i.e. the formula $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\mathsf{T}}$; hence:

$$\underline{\check{\boldsymbol{a}}} = \boldsymbol{S}^{-1}\underline{\boldsymbol{F}\check{\boldsymbol{x}}} + O(\|\check{\boldsymbol{x}}\|^2). \tag{11}$$

We obtain two insights. First, in this non-linear economy, a shock in the idiosyncratic characteristics of the l^{th} principal component of $\boldsymbol{M}$ passes through to the investments of that principal component multiplied a factor of s_l^{-1}.

Second, we can understand how idiosyncratic productivity shocks affects aggregate volatility. In fact,

$$\operatorname{Var}\left(\sum_i a_i\right) = \mathbb{E}\left[\check{\boldsymbol{a}}^{\mathsf{T}}\check{\boldsymbol{a}}\right] = \mathbb{E}\left[\underline{\check{\boldsymbol{a}}}^{\mathsf{T}}\underline{\check{\boldsymbol{a}}}\right] = \mathbb{E}\left[(\underline{\boldsymbol{F}\check{\boldsymbol{x}}})^{\mathsf{T}}\boldsymbol{S}^{-2}(\underline{\boldsymbol{F}\check{\boldsymbol{x}}})\right] + O(\mathbb{E}\|\check{\boldsymbol{x}}\|^4)$$

In the last line we plugged in (11). We have taken the terms that appear such as $\mathbb{E}[\check{\boldsymbol{a}}^{\mathsf{T}}\boldsymbol{v}]$, where $\boldsymbol{v} = O(\|\check{\boldsymbol{x}}\|_2)$, and bounded them by $O(\mathbb{E}\|\check{\boldsymbol{x}}\|^4)$. This follows because $\mathbb{E}[\check{\boldsymbol{a}}]$ is close to 0, with an error of order $O(\|\boldsymbol{x}\|^2)$.

Thus, our analysis of variance-reduction generalizes as long as the quadratic term in $\check{\boldsymbol{x}}$ is small enough to neglect. More precisely, if the planner can control variance of $\check{\boldsymbol{x}}$ subject to Assumption 2, the optimal variance reduction will be such that the variance reduction of lth principal component of $\boldsymbol{F}\check{\boldsymbol{x}}$ is decreasing in s_l.

8. Relation to other network measures

We discuss the relation between the key vectors that play a role in our characterizations—the principal components or singular vectors—and related network statistics.

First principal component and eigenvector centrality: We work, for ease of exposition, in the case of a symmetric $\boldsymbol{G}$. Suppose the matrix is nonnegative in each entry and irreducible (i.e., that the corresponding graph is connected). By the Perron-Frobenius Theorem, $\boldsymbol{u}^1(\boldsymbol{G})$ is entry-wise positive; indeed, this vector is the Perron vector of the matrix, also known as the vector of individuals' eigenvector centralities. Thus, for instance, Proposition 2 implies that, under strategic complementarities, interventions that aim to maximize aggregate utilities should change individuals' incentives in proportion to their eigenvector centralities.

It is worth comparing this result with results that highlight the importance of Bonacich centrality. Under strategic complements, equilibrium actions are proportional to the individuals' Bonacich centralities in the network $\boldsymbol{G}$ (Ballester et al., 2006). Within the Ballester et al. (2006) framework, it can easily be deduced that, if the objective of the planner is linear in the sum of actions then, under a quadratic cost function, the planner will target individuals proportionally to their Bonacich centralities (see also Demange (2017)). Bonacich centrality converges[39] to eigenvector centrality as the spectral radius of $\beta\boldsymbol{G}$ tends to 1, but, otherwise (e.g., for the β we used in Section 3) the two vectors can be quite different.

The substantive point is that the objective of our planner is to maximize aggregate equilibrium *utilities*, not actions, and that explains the difference in the targeting strategy. Indeed, our planner's objective can be written as follows (introducing a different constant factor for convenience):

$$\begin{aligned}\sum_i u_i \propto \frac{1}{n}\sum_i a_i^2 &= \left(\frac{1}{n}\sum_i a_i\right)^2 + \frac{1}{n}\sum_i \left(a_i - \frac{1}{n}\sum_i a_i\right)^2 \\ &= \bar{a}^2 + \sigma_{\boldsymbol{a}}^2,\end{aligned}$$

where $\sigma_{\boldsymbol{a}}^2$ is the variance of the action profile and $\overline{a}$ is the mean action. Thus, our planner likes to increase the sum of actions and also increases their diversity, simply as a mathematical consequence of his objective. This explains the reason his policies differ from those that would be in effect if just the mean action were the focus.

Last principal component: We have shown that, in games with strategic substitutes, for large budgets, interventions that aim to maximize aggregate utilities target agents proportionally to the eigenvector of $\boldsymbol{G}$ associated to the smallest eigenvalue of $\boldsymbol{G}$—the last principal component.

There is a connection between this result and the work of Bramoullé et al. (2014). Bramoullé et al. (2014) study the set of equilibria of a network game with linear best replies and strategic substitutes. They observe that such a game is a potential game, and derive explicitly the potential function. From this, they can deduce that the smallest eigenvalue of $\boldsymbol{G}$ is crucial for whether equilibrium is unique, and it is also useful for analyzing the stability of a particular equilibrium.[40] The basic intuition is that the magnitude of the smallest eigenvalue determines how small changes in individuals' actions propagate, via strategic substitutes, in the network. When these amplifications are strong, then multiple equilibria can emerge. Relatedly, when these amplifications are strong around an equilibrium, then that equilibrium will be unstable.

Our study of the strategic substitutes case is driven by different questions, and delivers different sorts of characterizations. We assume that there is a stable equilibrium, which is unique at least locally, and then we characterize optimal interventions in terms of the eigenvectors of $\boldsymbol{G}$. In general, all the eigenvectors—not just the one associated to the smallest eigenvalue—can matter. Interventions will focus *more* on the eigenvectors with smaller eigenvalues. When the budget is sufficiently large, the intervention will (in the setting of Section 5) focus only on the smallest-eigenvalue eigenvector. As we have discussed in Section 5.3, the network determinants of whether targeting is simple can be quite subtle. These considerations are all, to our knowledge, new in the study of network games.

[39]See, e.g., Calvó-Armengol et al. (2015) or Golub and Lever (2010).

[40]For stability of equilibrium what is relevant is the magnitude of the smallest eigenvalue of an appropriately defined subgraph of $\boldsymbol{G}$.

Nevertheless, at an intuitive level, there are important points of contact between our intuitions and those of Bramoullé et al. (2014). In our context, as we have discussed, our planner likes moving the incentives of adjacent individuals in opposite directions. The eigenvector associated to the smallest eigenvalue emerges as the one identifying the best way to do this at a given cost, and the eigenvalue itself measures how intensely the strategic effects amplify. This "amplification" property involves forces similar to those that make the smallest eigenvalue important to stability and uniqueness in Bramoullé et al. (2014).

Spectral approaches to volatility. Acemoglu et al. (2016) give a general analysis of which network statistics matter for volatility of network equilibria, and Baqaee and Farhi (2017) apply this analysis to certain macroeconomic questions about aggregate volatility. Though both papers note the importance of eigenvector centrality in (their analogues of) the case of strategic complements, their main focus is on how the *curvature* of best responses changes the volatility of an aggregate outcome, and what "second order" (curvature-related) network statistics are important. Our focus on the SVD to understand which first-order shocks are most amplified is distinct from the technical approaches they propose.

9. Concluding remarks

We solve a class of network intervention problems by exploiting singular value decompositions of matrices that capture strategic interactions. This approach allows us to consider different interaction structures as well as strategic interactions of different types. Our results, therefore, can speak to applications ranging from interventions in schools leveraging peer effects among pupils, interventions in oligopoly markets with differentiated products to combat surplus losses due to market power, and interventions in production networks to control aggregate volatility.

The quantities that come out as being important are the principal components of matrices of strategic interactions. These network statics are simple to describe and may be of interest to empirical researchers. The singular value decomposition (and its application to principal component analysis) has been a powerful and illuminating tool in applied mathematics and across many fields of economics. This suggests that a variety of further insights may be drawn from them, in the network context, using the characterizations of optimal interventions that we have established. We briefly elaborate on two further applications that are the focus of work in progress.

A class of standard intervention problems is one where the planner designs a budget-balanced tax/subsidy scheme in order to improve the economic outcome. For example, in the supply chain pricing game that we have studied, a planner could tax some suppliers, thereby increasing their marginal costs, and then use that tax revenue to subsidize the other suppliers. The planner will solve a similar problem to the one we have studied here, with the important difference that he will face a different constraint—namely, a budget-balance constraint. In ongoing work Galeotti et al. (2018) show that the decomposition of the network that we employed in this paper is useful in deriving the optimal taxation scheme and, in turn, in determining the welfare gains that can be achieved in supply chains through appropriate interventions.

Another possible application is the study of optimal interventions in a social learning model. Social learning studies based on the DeGroot (1974) model of opinion evolution—by DeMarzo, Vayanos, and Zwiebel (2003), Golub and Jackson (2010), and others—point out that, in the long run, an individual's influence on society's consensus belief is proportional to his eigenvector centrality. Furthermore, the speed of convergence of these opinions crucially depends on the segregation and homophily patterns in the population. It is natural to ask how an external planner would alter initial opinions to decrease aggregate miscoordination or disagreement costs incurred over the course of a learning process. Our techniques allow us to study this question, analyzing, for example, what patterns of initial disagreement are most consequential for welfare.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80, 1977–2016.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2016): "Networks, Shocks, and Systemic Risk," in *Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. Rogers, Oxford University Press.

Alon, N. and N. Kahale (1997): "A Spectral Technique for Coloring Random 3-Colorable Graphs," *SIAM Journal on Computing*, 26, 1733–1748.

Angeletos, G.-M. and J. La'O (2013): "Sentiments," *Econometrica*, 81, 739–779.

Angeletos, G.-M. and A. Pavan (2007): "Efficient Use of Information and Social Value of Information," *Econometrica*, 75, 1103–1142.

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006): "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 74, 1403–1417.

Banerjee, A., A. Chandrashekhar, E. Duflo, and M. O. Jackson (2016): "Gossip: Identifying Central Individuals in a Social Network," *Working Paper, MIT*.

Baqaee, D. R. and E. Farhi (2017): "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," Working paper, London School of Economics.

Belhaj, M. and F. Deroian (2017): "Contracting on Hetworks," Mimeo.

Bergemann, D., T. Heumann, and S. Morris (2015): "Information and Volatility," *The Journal of Economic Theory*, 158, 427–465.

——— (2017): "Information and Interaction," Cowles Foundation Discussion Paper 2088.

Bergemann, D. and S. Morris (2013): "Robust Predictions in Games with Incomplete Information," *Econometrica*, 81, 1251–1308.

Bloch, F. and N. Querou (2013): "Pricing in Social Networks," *Games and Economic Behavior*, 80, 263–281.

Borgatti, S. (2006): "Identifying Sets of Key Players in a Social Network," *Computational and Mathematical Organization Theory*, 12, 21–34.

Bramoulle, Y. and R. Kranton (2016): "Games Played on Networks," in *Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. Rogers, Oxford University Press.

Bramoullé, Y., R. Kranton, and M. d'Amours (2014): "Strategic Interaction and Networks," *The American Economic Review*, 104, 898–930.

Calvó-Armengol, A., J. Martí, and A. Prat (2015): "Communication and Influence," *Theoretical Economics*, 10, 649–690.

Candogan, O., K. Bimpikis, and A. Ozdaglar (2012): "Optimal Pricing in Networks with Externalities," *Operations Research*, 60, 883–905.

Chung, F. and M. Radcliffe (2011): "On the spectra of general random graphs," *the electronic journal of combinatorics*, 18, 215.

Condorelli, D. and A. Galeotti (2016): "Strategic Models of Intermediation Networks," in *Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. Rogers, Oxford University Press.

Dasaratha, K. (2017): "Distributions of Centrality on Networks," *CoRR*, abs/1709.10402.

de Martí, J. and Y. Zenou (2015): "Network games with incomplete information," *Journal of Mathematical Economics*, 61, 221–240.

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

Demange, G. (2017): "Optimal Targeting Strategies in a Network under Complementarities," *Games and Economic Behaviour*, 105, 84–103.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003): "Persuasion Bias, Social Influence, and Unidimensional Opinions," *The Quarterly Journal of Economics*, 118, 909–968.

Desai, M. and V. Rao (1994): "A Characterization of the Smallest Eigenvalue of a Graph," *Journal of Graph Theory*, 18, 181–194.

Dziubinski, M. and S. Goyal (2017): "How Do You Defend a Network?" *Theoretical Economics*, 12, 331–376.

Fainmesser, I. and A. Galeotti (2017): "Pricing Network Effects," *Review of Economic Studies*, 83, 165–198.

Feick, L. F. and L. L. Price (1987): "The Market Maven – A Diffuser of Marketplace Information," *Journal of Marketing*, 83–97.

Galeotti, A., B. Golub, S. Goyal, E. Talamàs, and O. Tamuz (2018): "Optimal Taxation in Supply Chains," *mimeo.*

Galeotti, A. and S. Goyal (2009): "Influencing the Influencers: A Theory of Strategic Diffusion," *The Rand Journal of Economics*, 40, 509–532.

Galeotti, A. and B. W. Rogers (2013): "Strategic Immunization and Group Structure," *American Economic Journal: Microeconomics*, 5, 1–32.

Golub, B. and M. O. Jackson (2010): "Naïve Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2, 112–49.

——— (2012): "How Homophily Affects the Speed of Learning and Best-Response Dynamics," *The Quarterly Journal of Economics*, 127, 1287–1338.

Golub, B. and C. Lever (2010): "The Leverage of Weak Ties: How Linking Groups Affects Inequality," Working paper, Harvard University.

Golub, B. and S. Morris (2017): "Expectations, Networks and Conventions," Available at SSRN: https://ssrn.com/abstract=2979086.

Goyal, S. (2017): "Markets and Networks," in *Advances in Economics: Eleventh World Congress of the Econometric Society*, ed. by M. P. B. Honore, A. Pakes and L. Samuelson, Cambridge University Press.

Goyal, S. and J. Moraga-Gonzalez (2001): "R&D Networks," *The Rand Journal of Economics*, 32, 686–707.

Hartfiel, D. and C. D. Meyer (1998): "On the Structure of Stochastic Matrices with a Subdominant Eigenvalue Near 1," *Linear Algebra and its Applications*, 272, 193–203.

Horn, R. A. and C. R. Johnson (2012): *Matrix Analysis*, New York, NY, USA: Cambridge University Press, 2nd ed.

Jackson, M. O. and Y. Zenou (2015): "Games on Networks," in *Handbook of the Game Theory with Economic Applications*, ed. by P. H. Young and S. Zamir, Elsevier.

Kempe, D., J. Kleinberg, and E. Tardos (2003): "Maximizing the Spread of Influence through a Social Network." in *Proceedings 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.*

Levin, D. A., Y. Peres, and E. L. Wilmer (2009): *Markov Chains and Mixing Times*, Providence, RI: American Mathematical Society.

Manea, M. (2016): "Models of Bilateral Trade in Networks," in *Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. Rogers, Oxford University Press.

Meyer, C. D. (2000): *Matrix Analysis and Applied Linear Algebra*, Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.

Rogers, E. (1983): *Diffusion of Innovations. Third Edition*, Free Press. New York.

Singh, N. and X. Vives (1984): "Price and Quantity Competition in a Differentiated Duopoly," *The Rand Journal of Economics*, 15, 546–554.

Urschel, J. C. (2018): “Nodal Decompositions of Graphs,” *Linear Algebra and its Applications*, 539, 60–71.

Valente, T. (2012): “Network Interventions,” *Science*, 337, 49–53.

Vives, X. (2001): *Oligopoly Pricing*, Cambridge, MA: The MIT Press.

Zenou, Y. (2016): “Key Players,” in *Oxford Handbook of the Economics of Networks*, ed. by Y. Bramoullé, A. Galeotti, and B. Rogers, Oxford University Press.

Appendix A. Omitted proofs

A.1. **Incentive-targeting.**

A.1.1. *Proof of Proposition 1.* The first step is to transform the maximization problem into the basis of the SVD, where it will be clearer which components should be optimally targeted. To this end, we first rewrite the cost and the objective in the SVD basis, using the fact that norms don't change under the orthogonal transformation $\boldsymbol{V}^{\mathsf{T}}$ which takes variables to their “underlined” coordinates:

$$K(\boldsymbol{b};\widehat{\boldsymbol{b}}) = \sum_i \left(b_i - \widehat{b}_i\right)^2 = \|\boldsymbol{b} - \underline{\boldsymbol{b}}\|_2^2 = \sum_{l=1}^n \left(\underline{b}_l - \underline{\widehat{b}}_l\right)^2,$$

and

$$\sum_{i\in\mathcal{N}} w_i^* = \frac{1}{2}\sum_{i\in\mathcal{N}} a_i^2 = \frac{1}{2}\|\boldsymbol{a}\|_2^2 = \frac{1}{2}\|\underline{\boldsymbol{a}}\|_2^2 = \frac{1}{2}\sum_{l=1}^n \underline{a}_l^2.$$

By defining

$$\alpha_l = \frac{1}{2(1-\beta\lambda_l(\boldsymbol{G}))^2}, \tag{12}$$

and, recalling (4), the maximization problem can be rewritten as

$$\max_{\underline{\boldsymbol{b}}} \sum_{l=1}^n \alpha_l \underline{b}_l^2 \qquad \text{(IT-SVD)}$$
$$\text{s.t. } \sum_{l=1}^n [\underline{b}_l - \underline{\widehat{b}}_l]^2 \le C.$$

We now transform the problem so that the control variable is $\underset{\sim}{\underline{\boldsymbol{b}}}$:

$$\max_{\underset{\sim}{\underline{\boldsymbol{b}}}} \sum_{l=1}^n \alpha_l \underline{\widehat{b}}_l^2 [\underset{\sim}{\underline{b}}_l + 1]^2 \qquad \text{(IT-SVD-REL)}$$
$$\text{s.t. } \sum_{l=1}^n \underline{\widehat{b}}_l^2 \underset{\sim}{\underline{b}}_l^2 \le C.$$

Note that, for all l the α_l defined by (12) are well-defined (by Assumption 1) and strictly positive. This has two implications.

First, if $\underset{\sim}{\boldsymbol{b}}^*$ solves (IT-SVD-REL), then the constraint in that problem binds. For otherwise, without violating the constraint in (IT-SVD-REL), we can slightly increase or decrease any $\underset{\sim}{b}_l^2$. Either the increase or the decrease is guaranteed to increase the corresponding $[\underset{\sim}{b}_l + 1]^2$ (since the α_l are all strictly positive).

Second, $\underset{\sim}{\boldsymbol{b}}^*$ satisfies $\underset{\sim}{b}_l^* \geq 0$ for every l. Suppose that for some l, we have $\underset{\sim}{b}_l^* < 0$. Then $[-\underset{\sim}{b}_l^* + 1]^2 > [\underset{\sim}{b}_l^* + 1]^2$. Since every α_l is positive, we can improve the objective without changing the cost by flipping the sign of $\underset{\sim}{b}_l^*$.

We now complete the proof by using the structure of the solution to (IT-SVD-REL) that follows from standard optimization theory. Observe that the Lagrangian corresponding to the maximization problem (IT-SVD-REL) is:

$$\mathcal{L} = \sum_{l=1}^{n} \alpha_l \widehat{\underline{b}}_l^2 \left[\underset{\sim}{b}_l + 1\right]^2 + \mu \left[C - \sum_{l=1}^{n} \widehat{\underline{b}}_l^2 \underset{\sim}{b}_l^2\right].$$

Taking our observation above that the constraint is binding at $\underset{\sim}{\boldsymbol{b}} = \underset{\sim}{\boldsymbol{b}}^*$ together with standard results on the Karush–Kuhn–Tucker conditions, the first-order conditions must hold exactly at the optimum with a positive μ:

$$0 = \frac{\partial \mathcal{L}}{\partial \underline{B}_l} = 2\widehat{\underline{b}}_l^2 \left[\alpha_l(1 + \underset{\sim}{b}_l^*) - \mu \underset{\sim}{b}_l^*\right] \quad l = 1, 2, \ldots, n. \tag{13}$$

We will take a generic $\widehat{\boldsymbol{b}}$ such that $\widehat{\underline{b}}_l \neq 0$ for each l. If for some l we had $\mu = \alpha_l$ then the right-hand side of (13) would be $2\widehat{\underline{b}}_l^2 \alpha_l$, which, by the generic assumption we just made and the positivity of α_l, would contradict (13). Thus the following holds with a nonzero denominator:

$$\underset{\sim}{b}_l^* = \frac{\alpha_l}{\mu - \alpha_l}. \tag{14}$$

It is immediate that if $\beta > 0$, α_l decreases in l and so $\underset{\sim}{b}_l^*$ decreases in l. If $\beta < 0$, α_l increases in l and so $\underset{\sim}{b}_l^*$ increases in l.

A.1.2. *Proof of Proposition 2:* [41]

Consider first the case $\beta > 0$. We begin by noting that, for $\boldsymbol{b}$ equal to the optimal solution,

$$\frac{\underset{\sim}{b}_l}{\underset{\sim}{b}_1} = \frac{\frac{\mu}{\alpha_1} - 1}{\frac{\mu}{\alpha_l} - 1} \leq \frac{\frac{\mu}{\alpha_1} - 1}{\frac{\alpha_1}{\alpha_l} - 1}$$

where the equality follows by (14) from the proof of Proposition 1, and the inequality follows because $\mu > \alpha_1$ (a fact also argued in the proof of Proposition 1). Now fix an $\eta > 0$. For any $l \neq 1$, a sufficient condition for $\underset{\sim}{b}_l / \underset{\sim}{b}_1 \leq \eta$ is that

$$\frac{\frac{\mu}{\alpha_1} - 1}{\frac{\alpha_1}{\alpha_2} - 1} < \eta, \text{ which holds if and only if } \frac{\mu}{\alpha_1} - 1 \leq \eta \left(\frac{\alpha_1}{\alpha_2} - 1\right) := \delta.$$

[41]We thank Fakhteh Saadatniaki for pointing out and correcting errors in a previous version of the proof.

This statement uses that $\alpha_l \leq \alpha_2$ for any such l.

Now, to complete the proof, the plan is as follows. In Part 1, we will show that if C exceeds a certain level, then $\frac{\mu}{\alpha_1} - 1 < \delta$. In Part 2, we will show that for a suitable η, the condition $\frac{\mu}{\alpha_1} - 1 < \delta$ implies a bound on the cosine similarity between $\Delta \boldsymbol{b}$ and $\boldsymbol{u}^1$ in terms of η. Finally, we will show that by setting η appropriately as a function of ϵ, these statements combine to deliver the result of the proposition.

Part 1. At the optimum, the constraint in (IT-SVD-REL) binds (as argued in the proof of Proposition 1), i.e.,

$$C = \sum_l \underline{\widehat{b}_l}^2 \left(\frac{\alpha_l}{\mu - \alpha_l} \right)^2 = \sum_l \underline{\widehat{b}_l}^2 \left(\frac{1}{\frac{\mu}{\alpha_l} - 1} \right)^2 .$$

Then note that if $\frac{\mu}{\alpha_1} - 1 \geq \delta > 0$ then $\left(\frac{1}{\frac{\mu}{\alpha_1} - 1} \right)^2 \leq \frac{1}{\delta^2}$ so

$$C = \sum_l \underline{\widehat{b}_l}^2 \left(\frac{1}{\frac{\mu}{\alpha_l} - 1} \right)^2 \leq \sum_l \underline{\widehat{b}_l}^2 \left(\frac{1}{\frac{\mu}{\alpha_1} - 1} \right)^2 \leq \frac{\|\underline{\widehat{\boldsymbol{b}}}\|_2^2}{\delta^2} = \frac{\|\widehat{\boldsymbol{b}}\|_2^2}{\delta^2},$$

where the first inequality follows because $\alpha_1 \geq \alpha_l$ and in the second inequality we have used the hypothesis that $\frac{\mu}{\alpha_1} - 1 \geq \delta$. Thus if $\frac{\mu}{\alpha_1} - 1 \geq \delta$ then $C \leq \frac{\|\widehat{\boldsymbol{b}}\|_2^2}{\delta^2}$.

Taking the contrapositive, if $C > \frac{\|\widehat{\boldsymbol{b}}\|_2^2}{\delta^2}$ then $\frac{\mu}{\alpha_1} - 1 < \delta$.

Part 2. Using the definition of δ in terms of η and recalling the definition of $\underset{\sim}{b}_k$, we have established that if $C > \frac{\|\widehat{\boldsymbol{b}}\|_2^2}{\delta^2}$, then

$$\frac{\underline{b}_l - \underline{\widehat{b}}_l}{\underline{b}_1 - \underline{\widehat{b}}_1} \frac{\underline{\widehat{b}}_1}{\underline{\widehat{b}}_l} = \frac{\underset{\sim}{b}_l}{\underset{\sim}{b}_1} < \eta.$$

It follows that

$$\frac{\underline{b}_l - \underline{\widehat{b}}_l}{\underline{b}_1 - \underline{\widehat{b}}_1} < \eta \cdot \frac{\underline{\widehat{b}}_l}{\underline{\widehat{b}}_1}.$$

Thus, summing across indices l not equal to 1,

$$\sum_{l \neq 1} (\Delta \underline{b}_l)^2 < (\Delta \underline{b}_1)^2 \cdot \frac{\eta^2}{\underline{\widehat{b}}_1^2} \sum_{l \neq 1} \underline{\widehat{b}}_l^2 .$$

Adding back in the $l = 1$ term,

$$\sum_l (\Delta \underline{b}_l)^2 < (\Delta \underline{b}_1)^2 \left[1 + \frac{\eta^2}{\underline{\widehat{b}}_1^2} \sum_{l \neq 1} \underline{\widehat{b}}_l^2 \right] .$$

From this we deduce

$$\|\Delta \underline{\boldsymbol{b}}\| < (\Delta \underline{b}_1) \left[1 + \frac{\eta^2}{\underline{\widehat{b}}_1^2} \sum_{l \neq 1} \underline{\widehat{b}}_l^2 \right]^{1/2}$$

or

$$\frac{\Delta \underline{b}_1}{\|\boldsymbol{\Delta} \underline{\boldsymbol{b}}\|} > \left[1 + \frac{\eta^2}{\widehat{\underline{b}}_1^2} \sum_{l \neq 1} \widehat{\underline{b}}_l^2\right]^{-1/2}.$$

Finally, using the fact that for any positive real number r, we have $[1+r]^{-1/2} \geq 1 - r/2$, we obtain

$$\frac{\Delta \underline{b}_1}{\|\boldsymbol{\Delta} \underline{\boldsymbol{b}}\|} > 1 - \frac{\|\widehat{\boldsymbol{b}}\|^2}{2\widehat{\underline{b}}_1^2} \eta^2.$$

Part 3. Now, if we set $\eta^2 = 2\epsilon \widehat{\underline{b}}_1^2 / \|\widehat{\boldsymbol{b}}\|^2$ the last statement becomes $\frac{\Delta \underline{b}_1}{\|\boldsymbol{\Delta} \underline{\boldsymbol{b}}\|} > 1 - \epsilon$. The condition on C that was sufficient to guarantee this was $C > \frac{\|\widehat{\boldsymbol{b}}\|_2^2}{\delta^2}$. Plugging in δ in terms of η and then η in terms of ϵ we find that this boils down to

$$C > \frac{\|\widehat{\boldsymbol{b}}\|_2^4}{2\widehat{\underline{b}}_1^2 \epsilon} \kappa_+.$$

The proof for the case of $\beta < 0$ is analogous, and therefore omitted.

A.1.3. *Proof of Corollary 1:* Consider $\beta > 0$. From Proposition 2, it is clear that as $C \to \infty$, we have $\underset{\sim}{b}_l^* / \underset{\sim}{b}_1^* \to 0$ for all $l \neq 1$. By definition of $\underset{\sim}{\boldsymbol{b}}$, this translates into

$$\frac{\underline{b}_l^* - \widehat{\underline{b}}_l}{\underline{b}_1^* - \widehat{\underline{b}}_1} \cdot \frac{\widehat{\underline{b}}_1}{\widehat{\underline{b}}_l} \to 0.$$

For the generic $\widehat{\boldsymbol{b}}$ we are considering, with all entries nonzero, the fraction $\widehat{\underline{b}}_1 / \widehat{\underline{b}}_l$ is a fixed constant. Thus for all $l \neq 1$, as $C \to \infty$,

$$\frac{\left[\underline{b}_l^* - \widehat{\underline{b}}_l\right]^2}{\left[\underline{b}_1^* - \widehat{\underline{b}}_1\right]^2} \to 0.$$

Since in (IT-SVD) in the proof of Proposition 1 (Section A.1.1) the constraint is

$$\sum_{l=1}^{n} \left[\underline{b}_l^* - \widehat{\underline{b}}_l\right]^2 = C.$$

Now, dividing the previous equation by $\left[\underline{b}_1^* - \widehat{\underline{b}}_1\right]^2$, and using the previous statement about the limit, makes it clear that $\left[\underline{b}_1^* - \widehat{\underline{b}}_1\right]^2 / C \to 1$, and thus, from the constraint, $\left[\underline{b}_l^* - \widehat{\underline{b}}_l\right]^2 / C \to 0$ for $l \neq 1$.

The argument for $\beta < 0$ is analogous.

A.2. **Details for Section 5.3.** In this section, which fills out the details omitted in Section 5.3, we first carry out the analysis of the random network with community structure for the case of strategic complements. We then provide the proof for the characterization of the network coefficients κ_+ and κ_- within the example.

The lower bound in the case of strategic complements. We derive here how the the network coefficient κ_+ depends on the structure of the random network with community structure for the case of strategic complements. Without loss of generality, assume for this case that $h_{\mathrm{H}} \geq h_{\mathrm{V}}$, and suppose $h_{\mathrm{H}} > 1/2$, so that there is some positive tendency to associate with own type in the East-West dimension.

Fact 3. If the conditions of Section 5.3 hold and $kp(k) \gg \log^2 k$ (i.e. the graph is dense enough) then, as $k \to \infty$,

$$\kappa_+ \xrightarrow{p} \left[\left(\frac{1-\beta(2h_{\mathrm{H}}-1)}{1-\beta} \right)^2 - 1 \right]^{-1}. \tag{15}$$

Hence, κ_+ becomes large as h_{H} approaches 1. Intuitively, in that limit, East and West are essentially segregated from each other and the targeting problem nearly decouples across those groups. Thus, there is no good one-dimensional summary of relative targeting and, therefore, the optimal intervention will depend on the initial conditions. For instance, if East happens to start out with higher initial $\widehat{b}_i$'s, then more effort will be focused on East, because the marginal returns are higher.[42]

Part 2. Let $\boldsymbol{P}$ be a 4-by-4 matrix that describes the probability of a link between agents in various groups. Defining the shorthand $\bar{h} = 1 - h$, the matrix $\boldsymbol{P}$ is as follows, with the $\circ$ entries filled in by symmetry:

$$\begin{array}{c c}
 & \begin{array}{cccc} \mathrm{NW} & \mathrm{NE} & \mathrm{SW} & \mathrm{SE} \end{array} \\
\begin{array}{c} \mathrm{NW} \\ \mathrm{NE} \\ \mathrm{SW} \\ \mathrm{SE} \end{array} &
\begin{bmatrix}
h_{\mathrm{NS}}h_{\mathrm{EW}} & h_{\mathrm{NS}}\bar{h}_{\mathrm{EW}} & \bar{h}_{\mathrm{NS}}h_{\mathrm{EW}} & \bar{h}_{\mathrm{NS}}\bar{h}_{\mathrm{EW}} \\
\circ & h_{\mathrm{NS}}h_{\mathrm{EW}} & \bar{h}_{\mathrm{NS}}\bar{h}_{\mathrm{EW}} & \bar{h}_{\mathrm{NS}}h_{\mathrm{EW}} \\
\circ & \circ & h_{\mathrm{NS}}h_{\mathrm{EW}} & h_{\mathrm{NS}}\bar{h}_{\mathrm{EW}} \\
\circ & \circ & \circ & h_{\mathrm{NS}}h_{\mathrm{EW}}
\end{bmatrix}
\end{array}$$

We will show $\boldsymbol{P}$ has four eigenvalues, and these are

$$1, \quad 2h_{\mathrm{NS}} - 1, \quad 2h_{\mathrm{EW}} - 1, \quad (2h_{\mathrm{NS}} - 1)(2h_{\mathrm{EW}} - 1). \tag{16}$$

The proof is to note that $\boldsymbol{P}$ can be viewed as the Kronecker product of two 2×2 matrices,

$$\boldsymbol{P} = \begin{bmatrix} h_{\mathrm{NS}} & \bar{h}_{\mathrm{NS}} \\ \bar{h}_{\mathrm{NS}} & h_{\mathrm{NS}} \end{bmatrix} \otimes \begin{bmatrix} h_{\mathrm{EW}} & \bar{h}_{\mathrm{EW}} \\ \bar{h}_{\mathrm{EW}} & h_{\mathrm{EW}} \end{bmatrix}.$$

The eigenvalues of the first matrix are 1 and $2h_{\mathrm{NS}} - 1$, and similarly for the second. By a standard fact, all eigenvalues of the Kronecker product are obtained by multiplying any eigenvalue of the first matrix by any eigenvalue of the second.

By standard results on the spectra of large random matrices (Chung and Radcliffe, 2011; Dasaratha, 2017), we can characterize the spectrum of $\boldsymbol{G}(k)$ as follows. Fix $\epsilon > 0$. Let

[42] When $\widehat{b}_i$ are higher, actions a_i will tend to be larger, and because of the complementarity between individuals' actions, it will be easier to increase any given individual's action in that group.

$d(k)$ be the average degree in graph k. First, by a standard Cheeger bound, under the degree-growth condition $kp(k) \gg \log^2 k$, with probability $1-\epsilon$, the degree $d_i(k)$ of any node has the property that $d_i(k)/d(k)$ is within ϵ of 1. Define $\widehat{\boldsymbol{G}}(k)$ by $\widehat{g}_{ij} = g_{ij}/\sqrt{d_i d_j}$. In the L^2 operator norm, this matrix is within ϵ of $\widetilde{\boldsymbol{G}}(k)$ defined by $\widetilde{g}_{ij} = g_{ij}/d(k)$, which is simply a scalar multiple of $\boldsymbol{G}(k)$. By the main result of Chung and Radcliffe (2011), the eigenvalues of $\widetilde{\boldsymbol{G}}$ outside $(-\epsilon, \epsilon)$ converge to the nonzero eigenvalues of $\boldsymbol{P}$ as $k \to \infty$, and, moreover, the two sets have the same cardinality. This shows that the largest eigenvalue of of $\boldsymbol{G}$ is approximately equal to its average degree, which justifies our claim about how we can select β. Now to establish the two facts, it suffices simply to compute the eigenvalue-dependent expressions from Section 5.2 for $\mathbf{P}$ and use the approximation result we have stated.

A.3. **Volatility reduction.**

Proof of Proposition 3. Take $\beta > 0$. Given our normalization $\mathbb{E}[\boldsymbol{b}] = 0$, the variance of aggregate investment, for any $\boldsymbol{\Omega}$, is

$$\begin{aligned}
\mathrm{Var}_{\boldsymbol{\Omega}}\left(\sum_i a_i\right) &= \mathbb{E}_{\boldsymbol{\Omega}}[\boldsymbol{a}^\mathsf{T}\boldsymbol{a}] \\
&= \sum_l \frac{\mathbb{E}_{\boldsymbol{\Omega}}[\underline{b}_l^2]}{(1-\beta\lambda_l)^2} \\
&= \sum_l \frac{\underline{\omega}_{ll}}{(1-\beta\lambda_l)^2},
\end{aligned}$$

where in the second line we have changed into the SVD basis by setting $\underline{\boldsymbol{\Omega}} = \boldsymbol{U}^\mathsf{T}\boldsymbol{\Omega}\boldsymbol{U}$.

Now applying this with $\boldsymbol{\Omega} = \boldsymbol{\Omega}^*$ and $\boldsymbol{\Omega} = \widehat{\boldsymbol{\Omega}}$ we find

$$\mathrm{Var}_{\boldsymbol{\Omega}^*}\left(\sum_i a_i\right) = \mathrm{Var}_{\widehat{\boldsymbol{\Omega}}}\left(\sum_i a_i\right) + \sum_l \frac{\underline{\omega}_{ll} - \widehat{\underline{\omega}}_{ll}}{(1-\beta\lambda_l)^2}.$$

The change in volatility of aggregate investment is a weighted sum of the decreases in the variance of each principal component, and the weight, $(1-\beta\lambda_l)^{-2}$, on the variance of principal component l of $\boldsymbol{G}$ is an increasing function of its eigenvalue λ_l. because $\beta > 0$.

Let $\boldsymbol{D} = \boldsymbol{\Omega}^* - \widehat{\boldsymbol{\Omega}}$ and $\underline{\boldsymbol{D}} = \underline{\boldsymbol{\Omega}}^* - \widehat{\underline{\boldsymbol{\Omega}}}$. (Note that $\underline{\Delta}_l$ defined in Proposition 3 is equal to $\underline{D}_{ll}$.) There is a permutation matrix (and therefore an orthogonal matrix) $\boldsymbol{P}$ so that $\widetilde{\underline{\boldsymbol{D}}} := \boldsymbol{P}\underline{\boldsymbol{D}}\boldsymbol{P}^\mathsf{T}$ has increasing entries along the diagonal. We can define

$$\widetilde{\boldsymbol{D}} = \boldsymbol{U}\widetilde{\underline{\boldsymbol{D}}}\boldsymbol{U}^\mathsf{T} = (\boldsymbol{U}\boldsymbol{P})\boldsymbol{D}(\boldsymbol{U}\boldsymbol{P})^\mathsf{T}.$$

By the assumption on K, the cost of the change

$$\widehat{\boldsymbol{\Omega}} \to \widehat{\boldsymbol{\Omega}} + \widetilde{\boldsymbol{D}}$$

is the same as the cost of the change

$$\widehat{\boldsymbol{\Omega}} \to \widehat{\boldsymbol{\Omega}} + \boldsymbol{D}.$$

In other words,

$$K(\widehat{\boldsymbol{\Omega}} + \widetilde{\boldsymbol{D}}; \widehat{\boldsymbol{\Omega}}) = K(\widehat{\boldsymbol{\Omega}} + \boldsymbol{D}; \widehat{\boldsymbol{\Omega}}).$$

But $\mathrm{Var}_{\widetilde{\boldsymbol{\Omega}^*}}\left(\sum_i a_i\right)$ is lower under $\widehat{\boldsymbol{\Omega}} + \widetilde{\boldsymbol{D}}$, unless the variance-covariance matrix did not change in this transformation, which could be the case if and only if the ordering of the reductions $\underline{\Delta}_l$ was already as the result claims. This proves the claim for $\beta > 0$; the proof for $\beta < 0$ is analogous and omitted. □

A.4. **Pricing game.**

Lemma 1. Consider the pricing game and consider a change from $\widehat{\boldsymbol{\Omega}}$ to $\boldsymbol{\Omega}$ the change in the expected consumer surplus, producer surplus and welfare are given by expressions 8-10.

Proof. Consumer surplus equals the equilibrium utility of the representative consumers, i.e., $CS(\boldsymbol{p}) = \frac{1}{2}\sum_f [Q_f(\boldsymbol{p})]^2$, and therefore

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\Omega}}\left[CS(\boldsymbol{p})\right] &= \frac{1}{2}\mathbb{E}_{\boldsymbol{\Omega}}\left[[\gamma\mathbf{1}_F - \boldsymbol{T}^\mathsf{T}\boldsymbol{p}]^\mathsf{T}[\gamma\mathbf{1}_F - \boldsymbol{T}^\mathsf{T}\boldsymbol{p}]\right] \\
&= \frac{1}{2}\mathbb{E}_{\boldsymbol{\Omega}}\left[\gamma^2\mathbf{1}_F^\mathsf{T}\mathbf{1} - 2\gamma\mathbf{1}_F^\mathsf{T}\boldsymbol{T}^\mathsf{T}\boldsymbol{p}\right] + \frac{1}{2}\mathbb{E}_{\boldsymbol{\Omega}}\left[\underline{\boldsymbol{p}}^\mathsf{T}\boldsymbol{\Lambda}\underline{\boldsymbol{p}}\right],
\end{aligned}$$

where in the last equation we have used the SVD of $\boldsymbol{T}\boldsymbol{T}^\mathsf{T}$. Hence

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\Omega}}\left[CS(\boldsymbol{p})\right] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}\left[CS(\boldsymbol{p})\right] &= \frac{1}{2}\mathbb{E}_{\boldsymbol{\Omega}}\left[\underline{\boldsymbol{p}}^\mathsf{T}\boldsymbol{\Lambda}\underline{\boldsymbol{p}}\right] - \frac{1}{2}\mathbb{E}_{\widehat{\boldsymbol{\Omega}}}\left[\underline{\boldsymbol{p}}^\mathsf{T}\boldsymbol{\Lambda}\underline{\boldsymbol{p}}\right] \\
&= \frac{1}{2}[\boldsymbol{I} + \boldsymbol{\Lambda}]^{-2}\boldsymbol{\Lambda}\left[\mathbb{E}_{\boldsymbol{\Omega}}\left[\underline{\boldsymbol{b}}^\mathsf{T}\underline{\boldsymbol{b}}\right] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}\left[\underline{\boldsymbol{b}}^\mathsf{T}\underline{\boldsymbol{b}}\right]\right] \\
&= \frac{1}{2}[\boldsymbol{I} + \boldsymbol{\Lambda}]^{-2}\boldsymbol{\Lambda}\left[\mathbb{E}_{\boldsymbol{\Omega}}\left[\underline{\boldsymbol{c}}^\mathsf{T}\underline{\boldsymbol{c}}\right] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}\left[\underline{\boldsymbol{c}}^\mathsf{T}\underline{\boldsymbol{c}}\right]\right],
\end{aligned}$$

which is equivalent to (8). Next, $PS(\boldsymbol{p}) = \sum_i \Pi_i(\boldsymbol{p})$ and so

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\Omega}}[PS(\boldsymbol{p})] &= \mathbb{E}\left[[\boldsymbol{p} - \boldsymbol{c}]^\mathsf{T}[\boldsymbol{p} - \boldsymbol{c}]\right] \\
&= \mathbb{E}_{\boldsymbol{\Omega}}\left[\underline{\boldsymbol{p}}^\mathsf{T}\underline{\boldsymbol{p}} + \underline{\boldsymbol{c}}^\mathsf{T}\underline{\boldsymbol{c}} - 2\underline{\boldsymbol{p}}^\mathsf{T}\underline{\boldsymbol{c}}\right]
\end{aligned}$$

Hence

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\Omega}}\left[PS(\boldsymbol{p})\right] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}\left[PS(\boldsymbol{p})\right] &= \mathbb{E}_{\boldsymbol{\Omega}}\left[\underline{\boldsymbol{p}}^\mathsf{T}\underline{\boldsymbol{p}} + \underline{\boldsymbol{c}}^\mathsf{T}\underline{\boldsymbol{c}} - 2\underline{\boldsymbol{p}}^\mathsf{T}\underline{\boldsymbol{c}}\right] \\
&= \left[(\boldsymbol{I} + \boldsymbol{\Lambda})^{-2} + \boldsymbol{I} - 2(\boldsymbol{I} + \boldsymbol{\Lambda})^{-1}\right]\left[\mathbb{E}_{\boldsymbol{\Omega}}[\underline{\boldsymbol{c}}^\mathsf{T}\underline{\boldsymbol{c}}] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}[\underline{\boldsymbol{c}}^\mathsf{T}\underline{\boldsymbol{c}}]\right] \\
&= \frac{\boldsymbol{\Lambda}^2}{(\boldsymbol{I} + \boldsymbol{\Lambda})^{-2}}\left[\mathbb{E}_{\boldsymbol{\Omega}}[\underline{\boldsymbol{c}}^\mathsf{T}\underline{\boldsymbol{c}}] - \mathbb{E}_{\widehat{\boldsymbol{\Omega}}}[\underline{\boldsymbol{c}}^\mathsf{T}\underline{\boldsymbol{c}}]\right],
\end{aligned}$$

which is equivalent to (9). The expression for (10) follows by combining (8) and (9). □

A.5. **Proof of Proposition 5.** The key point in the proof is to use our assumptions on K to ensure that it is an increasing function of $\|\boldsymbol{b} - \widehat{\boldsymbol{b}}\|_2$. Indeed, for any $\boldsymbol{x}$, we can find an orthogonal transformation $\boldsymbol{O}$ that maps $\boldsymbol{x}$ to $\|\boldsymbol{x}\|_2\mathbf{1}$. Thus, by Assumption 3, K is fully determined by its values on $K(\widehat{\boldsymbol{b}} + s\mathbf{1}; \widehat{\boldsymbol{b}})$ as s ranges over $[0, \infty]$. We know by Assumption 4 that these are increasing in s, let's say according to some function $k : \mathbb{R} \to \mathbb{R}$. Putting these facts together we can see that $K(\widehat{\boldsymbol{b}} + \boldsymbol{x}; \widehat{\boldsymbol{b}}) = k(\|\boldsymbol{x}\|_2)$.

Thus (IT-G) is equivalent to

$$\begin{aligned} &\max_{\boldsymbol{b}} \boldsymbol{a}^{\mathsf{T}}\boldsymbol{a} \\ &\text{s.t. } \|\boldsymbol{b} - \widehat{\boldsymbol{b}}\|_2^2 \leq C, \end{aligned} \tag{IT-G-S}$$

which is identical to the problem we studied in section 5 except for the lack of symmetry. The proof of proposition 5 is at this stage analogous to the proof of Proposition 1. The only difference is that we define $\underline{\boldsymbol{a}} = \boldsymbol{U}^{\mathsf{T}}\boldsymbol{a}$ and $\underline{\boldsymbol{b}} = \boldsymbol{V}^{\mathsf{T}}\boldsymbol{b}$, and define α_l in the proof of that result to be s_l, the corresponding singular value, avoiding the eigenvalues altogether.