

TARGETING INTERVENTIONS IN NETWORKS

ANDREA GALEOTTI, BENJAMIN GOLUB, AND SANJEEV GOYAL

ABSTRACT. Individuals interact strategically with their network neighbors, as in effort investment with spillovers among peers, or production decisions among firms connected by a supply chain. A planner can shape their incentives in pursuit of some goal—for instance, maximizing utilitarian welfare or minimizing the volatility of aggregate activity. We offer an approach to solving such intervention problems that exploits the singular value decomposition of network interaction matrices. The approach works by (i) describing the game in new coordinates given by the singular value decomposition of the network on which the game is played; and (ii) using that to deduce which components, and hence which individuals, a given type of intervention will focus on. Across a variety of intervention problems, simple orderings of the principal components characterize the planner’s priorities.

1. INTRODUCTION

Consider a group of individuals who interact strategically, with a network determining how a player’s action affects others’ incentives. An external entity—a *planner*—would like that the group attains a goal and seeks to achieve this goal through an intervention that changes individuals’ incentives. For instance, consider a group of school pupils who make choices about educational effort. Suppose that a pupil’s incentives to study are affected by the level of effort exerted by his friends. A utilitarian educational authority seeks to increase the sum of pupils’ utilities by offering subsidized tutoring or rewards for achievement to a few pupils. Whom should she target?

Alternatively, consider firms that are deciding how much to produce, with their decisions being strategically related—for instance, because they are involved in the same supply chains. The government can, at a cost, control the distributions of some fundamentals—for instance, by stabilizing the prices of some inputs. The government aims to maximize the expected

Date Printed. October 17, 2017.

We are grateful for helpful conversations with Xavier Vives and Glen Weyl. Gustavo Nicolas Paez and Eduard Talamás provided exceptional research assistance. Andrea Galeotti is grateful to the European Research Council for support through ERC-consolidator Grant Award. Galeotti: Department of Economics, European University Institute, and Department of Economics, University of Essex, Andrea.Galeotti@eui.eu. Golub: Department of Economics, Harvard University, bgolub@fas.harvard.edu; Goyal: Faculty of Economics and Christ’s College, University of Cambridge, sg472@cam.ac.uk.

consumer surplus. Which control efforts are most worthwhile, and how are they correlated across sectors?

This paper presents an approach to solving such intervention problems that relies on summarizing the structure of the network in its singular values and associated principal components and characterizing which ones a given type of intervention will focus on.

Consider now an abstract description of a network game, generalizing the situations we sketched above: individuals' payoffs depend on their own actions—which may be levels of effort, prices, or other quantities—and the actions of their contacts or neighbors. Individuals' actions may be strategic complements (pupils choosing levels of effort at studying, firms collaborating in joint research projects) or strategic substitutes (pricing decisions among competing firms in certain market structures, local public goods). The goal of the planner is—in our first problem—to maximize the sum of individuals' payoffs and she has at her disposal a limited amount of resources. The investment change agents' propensities to take the actions, and through the network effects this changes others' actions as well. We study the optimal allocation of these resources across individuals and how it depends on the strategic nature of interactions and the network describing who interacts with whom.

Let us start with strategic complements and suppose the planner uses her resources to subsidize individual efforts. A subsidy to an individual encourages this person to work harder. This raised effort, in turn, pushes up the efforts of his collaborators due to the strategic complementarity, which, in turn, increases the efforts of their collaborators, and so forth. Now, it is natural to suppose that the planner's marginal costs of raising any one individual's effort are increasing. This leads to the idea that the planner would like to target multiple individuals. So the planner is faced with the question of how best to coordinate her allocation of subsidy across the individuals. With strategic complements, she will want to move neighbors' incentives together, since increasing someone's effort makes it easier to increase the efforts of his neighbors'. The optimal policy will exploit this potential for amplification.

Next, consider the case of strategic substitutes. Targeting two neighbors and changing their incentives in the same direction would waste the planner's resources, because one individual's increase in effort would reduce his neighbors' incentives for effort; the different aspects of the planner's interventions would crowd each other out if she were to target neighbors. Thus, resources to increase incentives for activity should not be targeted at adjacent individuals.

Indeed, the best policy will often move neighbors' incentives in *opposite* directions to amplify the effects of an intervention and maximize social welfare.

Our goal is to generalize and formalize these intuitions, and understand how exactly network structure matters for determining the optimal targeting interventions. We show that the *singular value decomposition*—the decomposition of a linear system into orthogonal *principal components*—provides a simple and general formalization of these intuitions. In particular, in the case of strategic complements, a planner will focus on a “most representative component” of the network—the first principal component. This is a vector capturing a joint movement that maximizes amplification between neighbors. A subtler insight is that in the case of strategic substitutes, there is a different vector—the *last* principal component, corresponding to the smallest singular value—that is optimal. This targeting scheme minimizes crowd-out and leverages strategic substitutes.

The more general point is that the considerations involved in maximizing amplification of planners' efforts and avoiding crowd-out can be summed up by standard statistics of the network related to its singular value decomposition, and that this idea carries across a variety of applications. The mathematical reason behind this observation is that the singular value decomposition is very convenient for writing quadratic forms in the equilibrium actions. Quadratic forms turn up naturally in the analysis of the problems we have been discussing. In general, it is not only the first or last principal component that matters, but all the components of the decomposition. However, the degree to which they matter 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 economic activity in an investment setting, and two, maximizing consumer and producer surplus in a supply chain.

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 considering how policy interventions seeking to change outcomes can effectively exploit network structure, thereby economizing on scarce resources. Our paper contributes to a broad and exciting body of work—spread across economics, sociology, public health, marketing, and computer science, among other fields—which studies how to intervene in networks. In economics, recent work includes Ballester et al. (2006), Banerjee et al. (2016), Belhaj and Deroian (2017) Bloch and Querou (2013), Candogan et al. (2012), Demange (2017),

Dziubinski and Goyal (2017), Fainmesser and Galeotti (2017), Galeotti and Goyal (2009), Galeotti and Rogers (2013).¹

The novelty of the paper lies in (i) the formulation of a class of intervention problems; and (ii) in the proposal of a unified methodology to solve them. First, our planner’s problems share a general, economically natural, form—maximizing or minimizing an objective subject to a total constraint on resources spent to modify the environment. Within this class, the problems are rich in several ways: they involve a variety of planner’s objectives (e.g., total welfare, aggregate volatility); they involve different applications (educational effort, investment, pricing); and they consider different types of strategic interactions: strategic complements and substitutes. Second, across this class of problems, singular value decomposition—and the corresponding expression of a matrix via its principal components—offers a natural mathematical tool for identifying the optimal targets.

The rest of the paper is organized as follows. Section 2 presents the basic model, while Section 3 sets out notation and basic facts about the singular value decomposition and presents its application to a canonical network game. Section 4 solves for optimal targets when a planner can adjust individual incentives with the goal of maximizing social welfare. Section 5 presents and solves intervention problems in two other economic contexts. Section 6 shows how to relax restrictive assumptions used in the previous Sections. Appendix A contains 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 . Every individual has an exogenous *characteristic* $b_i \in \mathbb{R}$, with the vector of all characteristics denoted $\mathbf{b} \in \mathbb{R}^n$. Individual i chooses an action $a_i \in \mathbb{R}$, simultaneously with others; the vector of these is denoted $\mathbf{a} \in \mathbb{R}^n$. The payoffs to individual i given an action profile \mathbf{a} are:

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

So b_i denotes the individuals’ marginal benefits of her own action. The weighted, directed network with adjacency matrix \mathbf{G} has links (i, j) with weights g_{ij} ; it is a representation of strategic interactions.

¹Prominent contributions in related disciplines include Rogers (1983), Feick and Price (1987), Borgatti (2006), Kempe et al. (2003), and Valente (2012).

Suppose that \mathbf{G} is a nonnegative matrix. Then β captures the direction of strategic interdependencies: if $\beta > 0$, actions are strategic complements, and if $\beta < 0$, actions are strategic substitutes. We will not generally need the assumption of nonnegative \mathbf{G} , though the remarks we have just made show it can be a helpful case to think about. As Ballester et al. (2006) observed, any (pure strategy) Nash equilibrium action profile \mathbf{a} satisfies

$$[\mathbf{I} - \beta\mathbf{G}]\mathbf{a} = \mathbf{b}. \quad (1)$$

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

$$\mathbf{a} = [\mathbf{I} - \beta\mathbf{G}]^{-1}\mathbf{b}. \quad (2)$$

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

Assumption 1. The spectral radius of $\beta\mathbf{G}$ is less than 1.²

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

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

2.1. A basic network intervention problem. We now introduce a prototypical network intervention problem. The planner wishes to maximize aggregate utility of individuals; she has a budget and wishes to use this budget to modify the marginal benefits of individuals, for example by providing supplemental technologies to some of them. A *status quo* vector of characteristics $\widehat{\mathbf{b}}$ is given. The planner can change the vector $\widehat{\mathbf{b}}$ to \mathbf{b} , adjusting every individual's marginal benefit. The adjustment cost of changing $\widehat{\mathbf{b}}$ to \mathbf{b} is

$$K(\mathbf{b}; \widehat{\mathbf{b}}) = \sum_{i \in \mathcal{N}} (b_i - \widehat{b}_i)^2.$$

This reflects the idea that the planner faces increasing marginal costs as she seeks to make larger changes in individual's incentives.³

In terms of timing, the planner moves first with her intervention, and then the individuals engage in the simultaneous-move game to choose their actions. Assume that \mathbf{G} is such that

²Recall that the spectral radius of a matrix is the maximum absolute value of any of its eigenvalues.

³We extend the analysis to more general costs function in Section 6.

the equilibrium of the game is unique, and let w_i^* be the equilibrium utility of individual i given the planner's choice \mathbf{b} . Formally, then, the incentive-targeting problem is:

$$\begin{aligned} \max_{\mathbf{b}} \sum_{i \in \mathcal{N}} w_i^* & \quad (\text{IT}) \\ \text{s.t. } K(\mathbf{b}; \widehat{\mathbf{b}}) \leq C, & \end{aligned}$$

where C is the given resource constraint or budget.

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 both are costly. In particular, the planner may even choose to change the sign of b_i . In specific applications, it may be natural to assume that actions—such as research effort or prices—take on a positive value only. This could be reflected by appropriate constraints in the game; alternatively, we can study $\widehat{\mathbf{b}}$ large enough and C small enough to ensure that these constraints are respected at the optimal solution.

3. ANALYSIS OF THE GAME VIA THE SINGULAR VALUE DECOMPOSITION

This section defines notation for and recalls basic facts about the *singular value decomposition* (SVD) of a matrix. We then describe how it applies to the game, laying the groundwork for our main results.

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

$$\mathbf{M} = \mathbf{U}\mathbf{S}\mathbf{V}^\top, \quad (3)$$

where

- (1) \mathbf{S} is a $n \times m$ diagonal matrix with non-negative real numbers on the diagonal, s_l , called the *singular values* of \mathbf{M} ;
- (2) \mathbf{U} is an orthonormal $n \times n$ matrix whose columns are eigenvectors of $\mathbf{M}\mathbf{M}^\top$;
- (3) \mathbf{V} is an orthonormal $m \times m$ matrix whose columns are eigenvectors of $\mathbf{M}^\top\mathbf{M}$.

It is a standard fact that an SVD exists.⁴

Viewing \mathbf{M} as a map $M : \mathbb{R}^n \rightarrow \mathbb{R}^m$, the matrices \mathbf{V} and \mathbf{U} can be seen as bases for the domain and range, respectively, under which M is represented by a diagonal matrix. The

⁴A standard reference on SVD is Golub and Van Loan (1996).

columns of \mathbf{U} are called the *left-hand singular vectors* of \mathbf{M} , and the columns of \mathbf{V} are called the *right-hand singular vectors*. When we refer to *the l^{th} -ranked singular value* of a matrix \mathbf{M} , we mean the l^{th} -largest, and the l^{th} -ranked singular vector (on a given side) the corresponding column of \mathbf{U} or of \mathbf{V} .

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

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

An important application of the SVD is *principal component analysis*. We can think of the columns of \mathbf{M} as n data points. The first principal component of \mathbf{M} is defined as the n dimensional vector that minimizes the sum of squared distances to the actual \mathbf{M} . The first principal component is, therefore, a fictitious vector that best summarizes the data set \mathbf{M} . To characterize the other principal components, we orthogonally project all columns of \mathbf{M} off this vector and repeat this procedure. A well known result is that the left singular vectors of \mathbf{M} are, indeed, the principal components of \mathbf{M} ; a singular value quantifies the variation explained by the respective principal component. When we refer to the l^{th} principal component of \mathbf{M} we mean the l^{th} -ranked left singular vector of \mathbf{M} .

3.2. Analysis of the game using the SVD: the basic idea. Recall the equation characterizing the equilibrium of the game:

$$[\mathbf{I} - \beta \mathbf{G}] \mathbf{a} = \mathbf{b}.$$

The orthogonal decomposition entailed by the SVD gives a useful perspective on comparative statics and amplification of shocks in the system, which we will describe in this section. This idea is general and works for any \mathbf{G} and any β . To describe some implications simply, and to relate them to the structure of the network \mathbf{G} in a familiar way, we focus on a simple special case for exposition. Section 6 is devoted to stating the more general forms of the main results.

3.3. A special case. Assume that the matrix \mathbf{G} is *normal*, i.e. that $\mathbf{G}^{\top} \mathbf{G} = \mathbf{G} \mathbf{G}^{\top}$. This holds, for example, if \mathbf{G} is a symmetric matrix. The usefulness of this assumption is brought out in the following statement (see, e.g., Meyer (2000)).

Fact 1. If \mathbf{G} is normal and Assumption 1 holds, then there is a SVD of $\mathbf{M} = \mathbf{I} - \beta\mathbf{G}$ with $\mathbf{U} = \mathbf{V}$. This SVD can be chosen to correspond to a diagonalization $\mathbf{G} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$ so that the following hold:

1. $\mathbf{\Lambda}$ is a matrix of the eigenvalues of \mathbf{G} in decreasing order on the real line;
2. the i^{th} column of \mathbf{U} is the eigenvector of \mathbf{G} associated to the eigenvalue of \mathbf{G} in the position (i, i) of $\mathbf{\Lambda}$;
3. $\mathbf{S} = \mathbf{I} - \beta\mathbf{\Lambda}$.

Generically all the diagonal entries of $\mathbf{\Lambda}$, i.e., the eigenvalues of \mathbf{G} , are positive. The l^{th} eigenvector of \mathbf{G} , $\mathbf{u}^l(\mathbf{G})$, corresponds to the l^{th} principal component of \mathbf{G} . The decomposition is uniquely determined up to (i) a sign flip of any column of \mathbf{U} , and (ii) up to a permutation that reorders the eigenvalues in $\mathbf{\Lambda}$ and correspondingly reorders the columns of \mathbf{U} .

The implication of Fact 1 is that when \mathbf{G} is normal and Assumption 1 holds, the same orthogonal basis works on the left and the right in the SVD of $\mathbf{M} = \mathbf{I} - \beta\mathbf{G}$, and the relevant basis is one in which \mathbf{G} is diagonal. Furthermore, if $\beta < 0$,⁵ then the l^{th} principal component of $\mathbf{M} = \mathbf{I} - \beta\mathbf{G}$ is the l^{th} principal component of \mathbf{G} ; in the opposite case of $\beta > 0$, the l^{th} principal component of $\mathbf{M} = \mathbf{I} - \beta\mathbf{G}$ is the $(n - l + 1)^{\text{th}}$ principal component of \mathbf{G} .

3.4. Analysis of game using the SVD: details. In this subsection we discuss how the strategic structure of the game is illuminated by the SVD in the special case of a normal \mathbf{G} . For concrete examples of the principal components involved, see our application of the decomposition below in Section 4.2.

Substituting the expression $\mathbf{G} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$ into equation (1), we obtain

$$[\mathbf{I} - \beta\mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top]\mathbf{a} = \mathbf{b}$$

Multiplying the LHS and the RHS by \mathbf{U}^\top , we obtain an analogue of (2) characterizing the solution of the game:

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

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

$$a_l = \frac{1}{1 - \beta\lambda_l} b_l. \tag{4}$$

⁵Recall this corresponds to strategic substitutes if \mathbf{G} is nonnegative.

The equilibrium action in the l^{th} principal component of \mathbf{G} is simply a scaling of the magnitude of \mathbf{b} in that principal component.⁶ This is what we mean by a “decoupling” of the strategic interactions along the principal components: shocks to a given principal component are confined to that component in terms of their effect on actions.

In terms of magnitudes, suppose \mathbf{b} changes in the principal component of \mathbf{G} corresponding to a high value of $(1 - \beta\lambda_l)^{-1}$. In this case, the change in \mathbf{a} is large. With respect to the principal components of a (nonnegative) \mathbf{G} , when actions are strategic substitutes (resp. strategic complements), if the characteristic in the dimension of the l^{th} principal component of \mathbf{G} increases by ϵ , then the action of the l -th principal component goes up by less (resp. by more) than ϵ .

Rewriting in the original coordinates:

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

Thus the contributions to i 's action are 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 (b_l); and the magnification from the corresponding factor $(1 - \beta\lambda_l)^{-1}$. The SVD provides an useful way to understand the network locations that lead to higher or lower actions for a given attribute vector. Moreover, the SVD is very convenient for writing quadratic forms in the equilibrium actions, a fact that we exploit in the rest of the paper.

4. TARGETING INCENTIVES TO MAXIMIZE WELFARE

We are now in a position to state our first main result on optimal targeting of interventions for increasing utilitarian welfare in the problem (IT) of Section 2.1. Recall that the planner chooses the incentive vector \mathbf{b} and, with individuals playing the network game described in Section 2, has the following optimization problem:

$$\begin{aligned} \max_{\mathbf{b}} \quad & \sum_{i \in \mathcal{N}} w_i^* & \text{(IT)} \\ \text{s.t.} \quad & \sum_{i \in \mathcal{N}} (b_i - \widehat{b}_i)^2 \leq C, \end{aligned}$$

where the resource constraint is a given nonnegative number C , and $\widehat{\mathbf{b}}$ is a fixed vector of status quo attributes. We remark that the results we present in this Section generalize beyond the quadratic cost specification (see Proposition 5 in Section 6.2).

⁶Note that b_l is the magnitude of the orthogonal projection of \mathbf{b} onto column l of $\mathbf{V} = \mathbf{U}$.

4.1. **The structure of optimal interventions in terms of principal components.** Our first result shows that optimal interventions, in a suitable sense, focus on changing \mathbf{b} in some principal components more than others. Hence, the planner's priorities can be summarized in a general way, based on the eigenvalues of the network.

Recall that an arbitrary vector \mathbf{b} transformed into its SVD basis coordinates is denoted by $\underline{\mathbf{b}}$, with \underline{b}_l being the projection of \mathbf{b} onto the l th principal component of \mathbf{G} .⁷ It will be convenient for the statement of our result to work with *relative* changes in these components. Again for an arbitrary vector \mathbf{b} , let

$$\underline{b}_l = \frac{b_l - \widehat{b}_l}{\widehat{b}_l}$$

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

Proposition 1. Suppose \mathbf{G} is normal and Assumption 1 holds. Let \mathbf{b}^* be a solution to the incentive-targeting problem (IT) with graph \mathbf{G} . For a generic⁸ $\widehat{\mathbf{b}}$, we have that $\underline{b}_l^* > 0$ for all l , and that:

- (1) If $\beta > 0$ then \underline{b}_l^* is (weakly) decreasing in l ;
- (2) If $\beta < 0$ then \underline{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 an extreme principal component. This is the one corresponding to λ_1 , the largest eigenvalue of \mathbf{G} , or λ_n , the smallest eigenvalue of \mathbf{G} . Moreover, the degree of focus on principal components is monotonic in the eigenvalues. If $\beta > 0$ the degree of focus is decreasing in the order of principal components (ranked from greatest corresponding eigenvalue to least). On the other hand, when $\beta < 0$ the targeted budget is increasing in the order of the principal components. When \mathbf{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:

$$\max_{\underline{\mathbf{b}}} \sum_{l \in \mathcal{N}} \alpha_l \underline{b}_l^2 \tag{IT-SVD}$$

⁷We are fixing \mathbf{G} throughout this section, so that we may drop it as an argument on eigenvalues, etc.

⁸Nonzero in each component.

$$\text{s.t. } \sum_l [b_l - \widehat{b}_l]^2 \leq C.$$

This transformation uses that (i) orthonormal 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 \mathbf{G} is simply a scaling of the magnitude of $\underline{\mathbf{b}}$ in that principal component (recall (4)) by a coefficient we call α_l . In other words, the decoupling of strategic effects permits a convenient expression of the objective function. Then we make one more transformation, writing the objective and the constraint equivalently in terms of the *relative* changes, $\underline{\mathbf{b}}$:

$$\begin{aligned} \max_{\underline{\mathbf{b}}} \sum_{l \in \mathcal{N}} \alpha_l \widehat{b}_l^2 [b_l + 1]^2 & \quad \text{(IT-SVD-REL)} \\ \text{s.t. } \sum_l \widehat{b}_l^2 \underline{b}_l^2 \leq C. & \end{aligned}$$

From this it is straightforward to argue using basic optimization theory that at the optimal solution $\underline{\mathbf{b}}^*$, the entries \underline{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.

4.1.1. *The case of large budgets.* We next show that the optimal targeting strategy of the planner becomes extreme, focusing mostly on one component, when the resources available for intervention, C , are appropriately large. We begin with a quantitative result describing what happens to the planner’s focus as budgets become large.

Proposition 2. Fix a \mathbf{G} that is normal and generic⁹ such that Assumption 1 holds. Also fix a generic¹⁰ $\widehat{\mathbf{b}}$, and let \mathbf{b}^* be a solution to the incentive-targeting problem (IT).¹¹

1. Suppose $\beta > 0$. For any $\epsilon > 0$, if

$$C > \frac{\|\widehat{\mathbf{b}}\|_2^2}{\epsilon^2 \left[1 - \left(\frac{1-\beta\lambda_1}{1-\beta\lambda_2} \right)^2 \right]^2},$$

then $\underline{b}_l^*/\underline{b}_1^* < \epsilon$ for all $l \neq 1$.

⁹What is required is that λ_1 and λ_n are *strict* bounds on the other eigenvalues.

¹⁰Nonzero in each component.

¹¹This depends on the one parameter that has not yet been fixed, C .

2. Suppose $\beta < 0$. For any $\epsilon > 0$, if

$$C > \frac{\|\widehat{\mathbf{b}}\|_2^2}{\epsilon^2 \left[1 - \left(\frac{1 - \beta \lambda_n}{1 - \beta \lambda_{n-1}} \right)^2 \right]^2},$$

$$\underline{b}_l^* / \underline{b}_n^* < \epsilon, \text{ for all } l \neq n.$$

From this we can derive as an immediate corollary that most of the budget is spent on changing an extreme principal component (the one corresponding to $l = 1$ or $l = n$, depending on β), and the share of resources spent on changing the other components becomes negligible:

Corollary 1. Take the assumptions of Proposition 2:

- (1) Suppose $\beta > 0$. Then as $C \rightarrow \infty$, we have that $[b_l^* - \widehat{b}_l]^2 / C$ tends to 1 if $l = 1$, and to 0 otherwise.
- (2) Suppose $\beta < 0$. Then as $C \rightarrow \infty$, we have that $[b_l^* - \widehat{b}_l]^2 / C$ tends to 1 if $l = n$, and to 0 otherwise.

Corollary 1 captures the main take-away: large budgets C imply extreme focus. Proposition 2, the quantitative bound, describes how large C has to be, and what features of the network structure and the initial attributes $\widehat{\mathbf{b}}$ are important in ensuring that the focus is extreme. We now describe the content of the proposition in more detail. The initial attributes enter via $\|\widehat{\mathbf{b}}\|_2^2$; indeed, the proposition can be read as saying that $C / \|\widehat{\mathbf{b}}\|_2^2$ should exceed a number that depends only ϵ , β , and the network \mathbf{G} . This condition is harder to satisfy when the entries of $\widehat{\mathbf{b}}$ are larger or, holding the average of the entries fixed, when they are more variable.¹²

We now turn to the role the network plays in Proposition 2. When $\beta > 0$ (which corresponds to strategic complements assuming a nonnegative \mathbf{G}), the lower bound that $C / \|\widehat{\mathbf{b}}\|_2^2$ must exceed in Proposition 2 is $[1 - (1 - \beta \lambda_1)^2 / (1 - \beta \lambda_2)^2]^{-2}$; this is large when $\lambda_1 - \lambda_2$ is small. This quantity can be interpreted in terms of network structure. When $\lambda_1 - \lambda_2$ is small, the structural complexity of \mathbf{G} cannot be summarized by one principal component.¹³ Similarly, when actions are strategic substitutes, $\beta < 0$, heterogeneities in network locations are large when $\lambda_n - \lambda_{n-1}$ is large.

Corollary 1 implies that if actions are strategic complements, the optimal intervention \mathbf{b} is such that $\mathbf{b} - \widehat{\mathbf{b}}$ is (very nearly) proportional to the first principal component of \mathbf{G} , namely

¹²Recall that $\|\frac{1}{n}\widehat{\mathbf{b}}\|_2^2$ is equal to $\left(\frac{1}{n}\sum_i \widehat{b}_i\right)^2$ plus the variance of the entries of the vector $\widehat{\mathbf{b}}$.

¹³The number $\lambda_1 - \lambda_2$ is related to the *spectral gap*. See, e.g., Golub and Jackson (2012) for discussions of how the spectral gap corresponds to network structure, and in particular segregation and homophily.

$\mathbf{u}^1(\mathbf{G})$. On the other hand, if actions are strategic substitutes, the planner changes the individual incentives (very nearly) in proportion to the last principal component, $\mathbf{u}^n(\mathbf{G})$. Finally, when the initial attributes are zero ($\widehat{\mathbf{b}} = \mathbf{0}$), we can dispense with all the approximations. Assuming \mathbf{G} is generic in the sense used in Proposition 2, if $\widehat{\mathbf{b}} = \mathbf{0}$, then all of C is spent either (i) on increasing \underline{b}_1 (if $\beta > 0$), or (ii) on increasing \underline{b}_n (if $\beta < 0$).¹⁴

4.2. Examples. We now illustrate these results by considering two networks—a random network and a circle network—and two forms of strategic interaction—strategic complements and substitutes.¹⁵ In our computations, to bring out the effects of networks and interaction clearly, we take a large intervention budget, $C = 1000$. For the case of strategic complements we set $\beta = 0.1$, and for strategic substitutes we set $\beta = -0.1$. Figure 1 and Figure 2 present optimal interventions for the different treatments. The size of a node corresponds to the magnitude of the change made to that node’s attribute: i.e., $|b_i^* - \widehat{b}_i|$ at the optimal b_i^* . The colour reflects the direction of change: if node i is green (resp., red) it means that the intervention has increased (resp., decreased) the attribute from the initial value of \widehat{b}_i . Tables 1-4 present data on the initial b_i , the eigenvectors predicted to be important by the theory (the “first” and “last” eigenvectors), and on the optimal intervention ($\Delta b_i = b_i^* - \widehat{b}_i$), the change in action Δa_i , and the change in utility Δw_i .

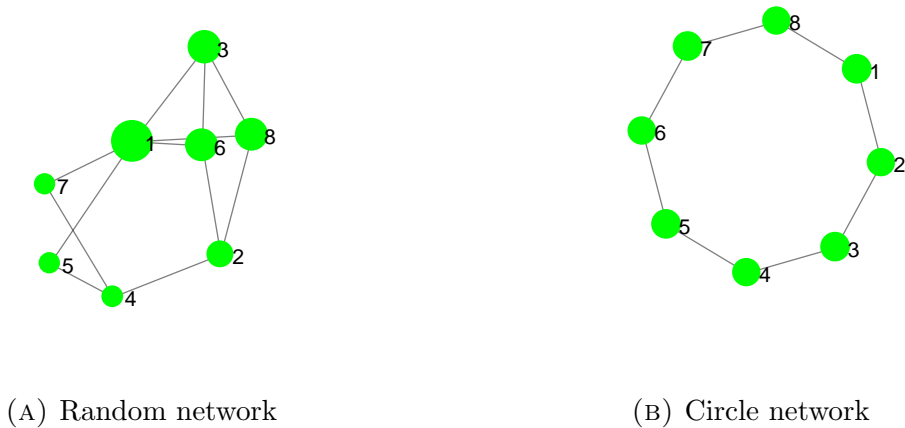


FIGURE 1. Optimal intervention with strategic complements

¹⁴Consider the form (IT-SVD) in our discussion after Proposition 1. Plug in $\widehat{\mathbf{b}} = \mathbf{0}$; then it is straightforward to show that if the focus is not monotonic, effort can be reallocated profitably among principal components without changing the cost.

¹⁵The random network is a Erdos-Renyi graph with $p = 0.5$, and it satisfies the spectral radius condition.

Node i	\hat{b}_i	\mathbf{u}_1 : “first” eigenvector	Δb_i	Δa_i	Δw_i
1	0.58	0.51	16.05	23.77	289.92
2	-0.16	0.32	9.92	14.68	110.59
3	0.38	0.40	12.64	18.63	182.49
4	-1.68	0.25	7.80	11.50	69.46
5	0.82	0.24	7.65	11.18	68.90
6	1.81	0.39	12.38	18.09	180.58
7	0.82	0.24	7.65	11.18	68.90
8	-0.19	0.39	12.37	18.08	180.25

TABLE 1. Targeting in a random network, $C = 1000$, $\beta = 0.1$

Node i	\hat{b}_i	\mathbf{u}_1 : “first” eigenvector	Δb_i	Δa_i	Δw_i
1	0.71	0.35	11.45	14.19	111.26
2	0.00	0.35	10.90	13.73	96.31
3	0.71	0.35	11.43	14.18	110.96
4	0.00	0.35	10.89	13.73	96.24
5	0.71	0.35	11.44	14.19	111.11
6	0.01	0.35	10.93	13.77	96.95
7	0.72	0.35	11.47	14.22	111.73
8	0.00	0.35	10.91	13.76	96.61

TABLE 2. Targeting in a circle network, $C = 1000$, $\beta = 0.1$

We start with strategic complements. Figure 1 compares optimal intervention in the random network and the circle network. Optimal intervention entails targeting the nodes in line with their entry in the eigenvector \mathbf{u}_1 , which is also called the *eigenvector centrality* (see Section 4.3). Thus, the optimal intervention raises the b_i of each node, and node i ’s eigenvector centrality determines the magnitude of this increase. In the random network, agent 1 has the highest eigenvector centrality and the change in his attribute is 16.05; by contrast, node 7 has the lowest eigenvector centrality and his attribute increases by only 7.65.

In the circle network nodes have the same structural positions. So any heterogeneity in targeting is due only to differences in \hat{b}_i .¹⁶ As Figure 2b illustrates, these initial differences are less important due to the large budget: the magnitude of the intervention and the consequent change in action is similar across nodes. So, the change in b_i ’s ranges from 10.9 to 11.5.

Figure 2 illustrates optimal intervention in the case of strategic substitutes. Now the intensity of intervention varies (roughly) in proportion to the “last” eigenvector, \mathbf{u}_n . This

¹⁶In both the networks we study, we choose a nonuniform initial vector $\hat{\mathbf{b}}$ for two reasons. First, it avoids nongeneric issues with some of the b_i being zero (as would be the case for a uniform $\hat{\mathbf{b}}$ in the circle). More importantly, it illustrates that the conclusions about targeting being in line with certain eigenvectors are not reliant on any particular structure of the $\hat{\mathbf{b}}$.

entails raising the b_i for some nodes and lowering the b_i of others. For example, 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 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. \square

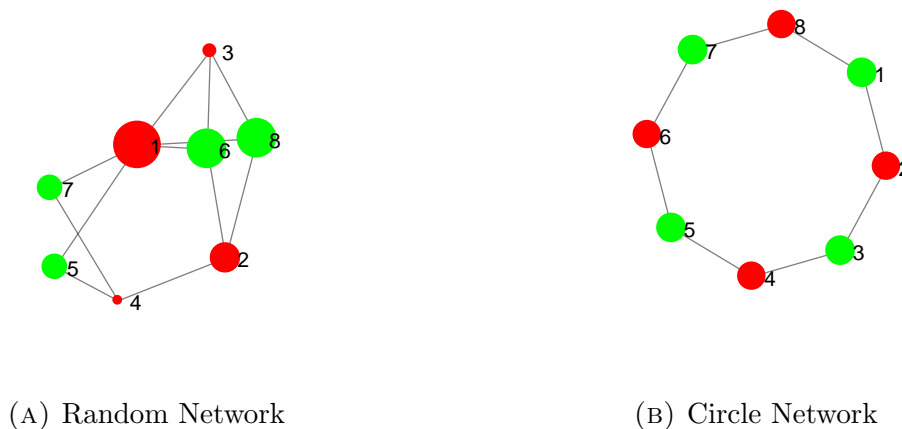


FIGURE 2. Optimal intervention with strategic substitutes

4.3. Relating principal components to network statistics and existing results. If \mathbf{G} is symmetric, its principal components are the eigenvectors of \mathbf{G} . In view of Propositions 1 and 2, as well as Corollary 1, we now discuss the first and the last eigenvectors.

First principal component and eigenvector centrality: Suppose our symmetric \mathbf{G} is nonnegative in each entry and irreducible (i.e., that the corresponding graph is connected). By the Perron-Frobenius Theorem, $\mathbf{u}^1(\mathbf{G})$ is entrywise positive; indeed, this vector is the Perron vector of the matrix, also known as the vector of individuals' *eigenvector centralities*. Thus, the eigenvector

Node i	\hat{b}_i	\mathbf{u}_n : “last” eigenvector	Δb_i	Δa_i	Δw_i
1	0.58	-0.55	-17.29	-22.7	265.29
2	-0.16	-0.34	-10.64	-13.98	100.59
3	0.38	-0.14	-4.32	-5.75	15.84
4	-1.68	-0.09	-2.73	-3.63	6.35
5	0.82	0.27	8.84	11.47	72.02
6	1.81	0.44	14.25	18.49	187.21
7	0.82	0.27	8.84	11.47	72.02
8	-0.19	0.44	14.23	18.48	186.77

TABLE 3. Targeting in a random network, $C = 1000$, $\beta = -0.1$

Node i	\hat{b}_i	\mathbf{u}_n : “last” eigenvector	Δb_i	Δa_i	Δw_i
1	0.71	0.35	11.46	14.21	111.57
2	0.00	-0.35	-10.92	-13.76	96.74
3	0.71	0.35	11.45	14.2	111.25
4	0.00	-0.35	-10.9	-13.74	96.43
5	0.71	0.35	11.42	14.17	110.8
6	0.01	-0.35	-10.89	-13.73	96.11
7	0.72	0.35	11.45	14.2	111.44
8	0.00	-0.35	-10.92	-13.77	96.84

TABLE 4. Targeting in a circle network, $C = 1000$, $\beta = -0.1$

centrality captures the distribution that is most representative of the various individuals’ neighborhoods. We have shown that, under strategic complementarities, interventions that aim to maximize aggregate utilities should mostly be focused on the eigenvector centrality.¹⁷

It is worth considering this result in light of another widely studied centrality statistic. Under strategic complements, equilibrium actions are proportional to the individuals’ Bonacich centralities in the network \mathbf{G} (Ballester et al., 2006). 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 to eigenvector centrality as the spectral radius of $\beta\mathbf{G}$ tends to 1, but, otherwise (and in particular for the β we have studied) the two vectors can be very 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

¹⁷The work on social learning by DeMarzo, Vayanos, and Zwiebel (2003) and Golub and Jackson (2010), based on the DeGroot (1974) model of opinions, draws attention to eigenvector centrality. They point out that, in the long run, an individual’s influence on society’s consensus belief is proportional to his eigenvector centrality.

planners objective can be written (we introduce a factor of $1/n$ for convenience) as:

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

where $\sigma_{\mathbf{a}}^2$ is the variance of the action profiles and \bar{a} is their mean. Thus, our planner likes to increase the sum of actions and also increases their diversity (not intrinsically, just as a mathematical consequence of his objective).

Last Principal Component: Comparative statics results for network games with strategic substitutes are less developed, and so the study of optimal intervention has been limited so far. All our results for games of strategic complements have analogues for the case of strategic substitutes. The SVD approach permits a simple characterization of optimal targets in the substitute case: the planner should focus, primarily, on the *last* principal component of \mathbf{G} . This principal component is the most idiosyncratic or least representative—what is left over after we have orthogonally projected, one by one, off all the other “more representative” components (see the discussion in Section 3.1). This is a new finding that has no analogues, to our knowledge, in the prior network games literature.¹⁸

5. TWO FURTHER APPLICATIONS

This section presents two additional network intervention problems: (i) minimizing aggregate volatility in investment and (ii) maximizing consumers and producers surplus in a supply chain. In both cases, we apply SVD to identify optimal targets.

5.1. Aggregate volatility. The network game we have presented is closely 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); see Acemoglu, Ozdaglar, and Tahbaz-Salehi (2016) for a survey of this literature.¹⁹ Other work has studied

¹⁸Bramoullé et al. (2014) provides conditions on the most negative eigenvalue of $\mathbf{I} - \beta\mathbf{G}$ in order for the network game to have a unique equilibrium. This eigenvalue is also informative above the stability of equilibrium. The *most negative* eigenvalue of $\mathbf{I} - \beta\mathbf{G}$ is very different from the eigenvalue of least magnitude of \mathbf{G} . Furthermore, the corresponding eigenvector is not studied in that paper, whereas it is key to the structure of optimal interventions for us.

¹⁹In the production network application, firms are price-taker firms, \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 is product j for the production of good i . See Acemoglu et al. (2016) for a formal connection to the network game presented here.

the role of private information in volatility, e.g., Angeletos and Pavan (2007) and Bergemann, Heumann, and Morris (2015), recently bringing in a network dimension (de Martí and Zenou, 2015; Bergemann, Heumann, and Morris, 2017).²⁰

Following this latter strand, we interpret actions as levels of investment. The vector \mathbf{b} of exogenous attributes is common knowledge among the individuals, but random from the perspective of the planner at the time of her decision. The variance-covariance matrix of the shock vector at the status quo is $\widehat{\Omega}$. The planner wishes to minimize the volatility in the level of aggregate investment:

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

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

Assumption 2. K is invariant to rotations of coordinates:

$$K(\widehat{\Omega} + \mathbf{D}; \widehat{\Omega}) = K(\widehat{\Omega} + \mathbf{O}^\top \mathbf{D} \mathbf{O}; \widehat{\Omega})$$

For example, suppose that the cost $K(\Omega; \widehat{\Omega})$ is equal to the reduction in the sum of attribute variances. Then $K(\Omega; \widehat{\Omega}) = \sum_i \widehat{\omega}_{ii} - \sum_i \omega_{ii} = \text{trace}(\widehat{\Omega} - \Omega)$. This specification turns out to satisfy Assumption 2.²¹

Under this assumption, we will study the variance-minimization problem described above—formally,

$$\begin{aligned} \min_{\Omega} \text{Var} \left(\sum_i a_i \right) & \quad \text{(VM)} \\ \text{s.t. } K(\Omega; \widehat{\Omega}) & \leq C. \end{aligned}$$

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

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

²⁰The focus of those papers is on understanding how private information shapes aggregate volatility in a game with linear best replies. Thus, 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.

²¹To see this, note that the trace is the sum of the eigenvalues, and this does not change under conjugacy transformations of the argument.

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 per principal component of \mathbf{G} is monotone in their ranking. In particular, the first principal component will receive the most focus when investments are strategic complements and the last principal component will receive the most focus when investments are strategic substitutes.

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, where the formula for $\text{Var}_{\Omega}(\sum_i a_i)$ is simple due to the ‘‘decoupling’’ of the system. It shows that the contribution to aggregate volatility of the variance of each principal component of \mathbf{G} 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 this rearrangement is feasible.

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

$$\begin{aligned} \text{Var}_{\Omega} \left(\sum_i a_i \right) &= \mathbb{E}_{\Omega}[\mathbf{a}^{\top} \mathbf{a}] \\ &= \sum_l \frac{\mathbb{E}_{\Omega}[b_l^2]}{(1 - \beta\lambda_l)^2} \\ &= \sum_l \frac{\underline{\omega}_l}{(1 - \beta\lambda_l)^2}, \end{aligned}$$

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

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

$$\text{Var}_{\Omega^*} \left(\sum_i a_i \right) = \text{Var}_{\hat{\Omega}} \left(\sum_i a_i \right) + \sum_l \frac{\underline{\omega}_l - \hat{\omega}_l}{(1 - \beta\lambda_l)^2}.$$

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

Let $\mathbf{D} = \Omega^* - \hat{\Omega}$ and $\underline{\mathbf{D}} = \underline{\Omega}^* - \underline{\hat{\Omega}}$. (Note that $\underline{\Delta}_l$ defined in Proposition 3 is equal to \underline{D}_l .) There is a permutation matrix (and therefore an orthonormal matrix) \mathbf{P} so that

$\tilde{\underline{D}} := \underline{P}\underline{D}\underline{P}^\top$ has decreasing entries along the diagonal. We can define

$$\tilde{\underline{D}} = \underline{U}\tilde{\underline{D}}\underline{U}^\top = (\underline{U}\underline{P})\underline{D}(\underline{U}\underline{P})^\top.$$

By the assumption on K , the cost of the change

$$\hat{\underline{\Omega}} \rightarrow \hat{\underline{\Omega}} + \tilde{\underline{D}}$$

is the same as the cost of the change

$$\hat{\underline{\Omega}} \rightarrow \hat{\underline{\Omega}} + \underline{D}.$$

In other words,

$$K(\hat{\underline{\Omega}} + \tilde{\underline{D}}; \hat{\underline{\Omega}}) = K(\hat{\underline{\Omega}} + \underline{D}; \hat{\underline{\Omega}}).$$

But $\text{Var}_{\hat{\underline{\Omega}}^*}(\sum_i a_i)$ is lower under $\hat{\underline{\Omega}} + \tilde{\underline{D}}$, 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 $\underline{\Delta}_i$ was already as the result claims. This proves the claim for $\beta > 0$; the proof for $\beta < 0$ is analogous and omitted. \square

5.2. Pricing in a supply chain. Next, we consider a pricing game between suppliers in a supply chain. The intervention reduces 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* $\mathcal{F} = \{1, 2, \dots, F\}$. Final goods are made using the set of *inputs* $\mathcal{N} = \{1, 2, \dots, N\}$; supplier i produces input $i \in \mathcal{N}$. Following Vives (2001); Singh and Vives (1984), a representative consumer with quadratic utilities chooses how much to consume of each final good. Given price vector $\underline{P} = \{P_1, \dots, P_F\}$, the utility of the consumer is

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

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 of final goods: $Q_f = \gamma - P_f$. The

utility of the representative consumer is

$$U^* = \frac{1}{2}Q_f^2.$$

We now describe how inputs are transformed into final goods. Let \mathbf{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 $i \in \mathcal{N}$, $\mathbf{t}_i \cdot \mathbf{t}_i = 1$.²² We assume that the final goods are competitive and so the price of final good f equals the marginal cost of production of good f .²³ We can thus write $P_f(\mathbf{p}) := \sum_{i \in \mathcal{N}} t_{if} p_i$, or, in matrix notation, $P(\mathbf{p}) = \mathbf{T}^\top \mathbf{p}$. The vector of demand for inputs is $q(\mathbf{p}) = \mathbf{T}Q(\mathbf{T}^\top \mathbf{p})$. The demand of supplier i , which depends on all prices, is

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

For a given price profile, \mathbf{p} , the profit of supplier i with a constant marginal cost c_i is

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

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 \mathbf{p} solves the system

$$[\mathbf{I} + \mathbf{T}\mathbf{T}^\top]\mathbf{p} = \mathbf{b}, \tag{5}$$

where $\mathbf{b} = \mathbf{c} + \gamma\mathbf{T}\mathbf{1}$. This is equivalent to system (1) with $\mathbf{G} = \mathbf{T}\mathbf{T}^\top$, $\beta = -1$ and the endogenous variables are suppliers' prices, $\mathbf{a} = \mathbf{p}$. 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 $\mathbf{G} = \mathbf{T}\mathbf{T}^\top$ is symmetric (and therefore normal); furthermore $\beta = -1$ and so Assumption 1 holds. Second, the SVD of $\mathbf{G} = \mathbf{T}\mathbf{T}^\top$ is related to the SVD of \mathbf{T} . Since \mathbf{T} may not be a square matrix, the SVD of \mathbf{T} reads $\mathbf{T} = \mathbf{U}\mathbf{S}\mathbf{V}^\top$, where the columns of \mathbf{V} and \mathbf{U} are the right and left singular vectors of \mathbf{T} , respectively. It follows that the SVD of $\mathbf{G} = \mathbf{T}\mathbf{T}^\top$ is given by $(\mathbf{U}, \mathbf{\Lambda}, \mathbf{U}^\top)$, where $\mathbf{\Lambda} = \mathbf{S}\mathbf{S}^\top = \mathbf{S}^2$. Hence, the singular values of \mathbf{G} are the square of the singular values of \mathbf{T} , and the principal components of \mathbf{G} , i.e., the right singular vectors of \mathbf{T} , are bundles of final goods, which are the best fit of the underlying technology of production \mathbf{T} .

²²This is a normalization: we choose the relevant units of each input i such that the Euclidean length of each vector \mathbf{t}_i is equal to one.

²³Constant markups can be added without significantly changing our analysis.

In this application, we refer to the columns of \mathbf{U} as fundamental bundles of final goods and we define $\underline{\mathbf{p}} = \mathbf{U}^\top \mathbf{p}$ and $\underline{\mathbf{b}} = \mathbf{U}^\top \mathbf{b}$. We can then rewrite the equilibrium price system (5) as follows:

$$\underline{\mathbf{p}} = [\mathbf{I} + \mathbf{\Lambda}]^{-1} \underline{\mathbf{b}} \iff p_l = \frac{b_l}{1 + \lambda_l}. \quad (6)$$

Now suppose that the production technology is common knowledge among market participants (the vector \mathbf{c} of marginal costs is common knowledge), 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 $\widehat{\mathbf{\Omega}}$ and the planner can change it to $\mathbf{\Omega}$ at a cost $K(\mathbf{\Omega} - \widehat{\mathbf{\Omega}})$, which satisfies Assumption 2.

We study the optimal choice of $\mathbf{\Omega}$ under the constraint $K(\mathbf{\Omega} - \widehat{\mathbf{\Omega}}) \leq C$, for three objectives: expected consumer surplus, $CS(\mathbf{p})$, producer surplus, $PS(\mathbf{p})$ and social welfare, $SW(\mathbf{p}) = CS(\mathbf{p}) + PS(\mathbf{p})$. The change of these three quantities, when we move from $\widehat{\mathbf{\Omega}}$ to $\mathbf{\Omega}$, turns out to be a convex combination of the respective changes in the variances of the marginal costs of the fundamental bundles of final goods, and the weight associated to the l^{th} -ranked fundamental bundle is a function of the l^{th} -ranked singular value. Formally,

$$\mathbb{E}_{\mathbf{\Omega}}[CS(\mathbf{p})] - \mathbb{E}_{\widehat{\mathbf{\Omega}}}[CS(\mathbf{p})] = \sum_l \frac{\lambda_l}{(1 + \lambda_l)^2} [\underline{\omega}_{ll} - \widehat{\underline{\omega}}_{ll}] \quad (7)$$

$$\mathbb{E}_{\mathbf{\Omega}}[PS(\mathbf{p})] - \mathbb{E}_{\widehat{\mathbf{\Omega}}}[PS(\mathbf{p})] = \sum_l \frac{\lambda_l^2}{(1 + \lambda_l)^2} [\underline{\omega}_{ll} - \widehat{\underline{\omega}}_{ll}] \quad (8)$$

$$\mathbb{E}_{\mathbf{\Omega}}[SW(\mathbf{p})] - \mathbb{E}_{\widehat{\mathbf{\Omega}}}[SW(\mathbf{p})] = \sum_l \frac{\lambda_l}{(1 + \lambda_l)} [\underline{\omega}_{ll} - \widehat{\underline{\omega}}_{ll}]. \quad (9)$$

where we recall from Section 5.1 that $\underline{\omega}_{ll} = \text{Var}(\mathbf{u}^l \cdot \mathbf{c})$ under $\mathbf{\Omega}$ and $\widehat{\underline{\omega}}_{ll} = \text{Var}(\mathbf{u}^l \cdot \mathbf{c})$ under $\widehat{\mathbf{\Omega}}$.

Proposition 4. Assume K satisfies Assumption (2).

1. Suppose $\mathbf{\Omega}^*$ maximizes expected producer surplus or expected total welfare. Then the variance reduction in the marginal cost of the l^{th} fundamental bundle of final goods is decreasing in l .
2. Suppose $\mathbf{\Omega}^*$ maximizes expected consumer surplus. Let the \bar{l}^{th} fundamental bundle of final goods be such that $\lambda(\bar{l}) \geq 1$ and $\lambda(\bar{l} + 1) < 1$.²⁴ The variance reduction in the marginal cost of the l^{th} fundamental bundle is increasing in l if $l \leq \bar{l}$ and, otherwise, decreasing.

²⁴Set $\bar{l} = n + 1$ if $\lambda(n) > 1$ and set $\bar{l}^{\text{th}} = 0$ if $\lambda(1) < 1$.

Proof of Proposition 4. Consider expression (8) and note that the weight to the variance reduction 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 (7) for the expected consumer surplus. Note that the weight to the variance reduction 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 (7)-(9) are derived in Lemma 1 in the Appendix. \square

There are two main effects identified by Proposition 4. The first effect is a *pass-through effect* across suppliers. The pricing game is a game of strategic substitutes and, therefore, shocks in marginal costs which alter the price of some suppliers are attenuated by the strategic response of other suppliers. This effect is summarized in expression (6) that indicates that shocks are attenuated the most along the highest-ranked 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

$$\mathbf{P} = \mathbf{T}^\top \mathbf{p} \iff \underline{\mathbf{P}} = \mathbf{S}^\top \underline{\mathbf{p}},$$

where the equivalence follows using the SVD of \mathbf{T} . Hence, the quantity effects are the stronger along the main principal components.²⁵

When the objective of the planner is to maximize consumer surplus, these two effects are countervailing: the pass-through effect pushes the planner to target 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 reduction 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 reduction is on the main principal components.

6. GENERALIZATIONS: NON-NORMAL \mathbf{G} , DIFFERENT COST FUNCTIONS, AND NONLINEAR SYSTEMS

The assumption of normal \mathbf{G} and linear systems has been convenient for presenting key aspects of our approach. However, these assumptions will often not hold in practical

²⁵If 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 $s_{l+1} \epsilon$

applications. To conclude the formal analysis, we show that many of the insights of the previous sections generalize to a case without these assumptions. 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 normal \mathbf{G} ; we then study a nonlinear system.

6.1. Incentive targeting: Beyond normal \mathbf{G} . We extend the analysis of Section 4 to non-normal \mathbf{G} . We describe this case within the following, slightly more general, framework. Individuals take endogenous actions $\mathbf{a} = (a_1, \dots, a_n)$ which, in equilibrium, satisfy

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

where $\widehat{\mathbf{b}}$ is a vector of idiosyncratic characteristic. In the setting of Section 4, we have $\mathbf{M} = \mathbf{I} - \beta\mathbf{G}$ for an arbitrary matrix \mathbf{G} such that \mathbf{M} is invertible.²⁶ The planner can intervene and change $\widehat{\mathbf{b}}$ to \mathbf{b} at a cost with the aim of maximizing the objective

$$W(\mathbf{a}) = F(\mathbf{a}^\top \mathbf{a}),$$

where F is an increasing function (at least on the domain of \mathbf{a} achievable through intervention). When F is a linear function the objective corresponds to the one studied in Section 4.

We use the SVD of \mathbf{M} to rewrite the argument of the objective function as a convex combination of the contribution that each principal component of \mathbf{M} has in determining the argument of F , and thus in turn 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 \mathbf{M} corresponds to the formula $\mathbf{M} = \mathbf{U}\mathbf{S}\mathbf{V}^\top$ and therefore, using the transformed coordinates $\underline{\mathbf{b}} = \mathbf{U}^\top \mathbf{b}$, we obtain

$$\mathbf{a}^\top \mathbf{a} = \underline{\mathbf{b}}^\top [\mathbf{S}^\top \mathbf{S}]^{-1} \underline{\mathbf{b}} = \sum_{l=1}^n \frac{1}{s_l^2} b_l^2.$$

Hence, the objective function is

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

It is now apparent that the analysis of optimal intervention can be carried out in the same way as in Section 4. The result is that, under the optimal intervention, the degree of focus of

²⁶This amounts to the (generically satisfied) requirement that $1/\beta$ is not an eigenvalue of \mathbf{G} .

the planner to principal component l of \mathbf{M} is decreasing in its corresponding singular value s_l .²⁷

6.2. More general cost functions in incentive-targeting. Consider any function $(\mathbf{b}; \widehat{\mathbf{b}}) \mapsto K(\mathbf{b}; \widehat{\mathbf{b}})$ and assume the following about it.

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

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

That is, the costs of interventions are rotationally symmetric, in the sense that they stay fixed when we rotate a given change vector \mathbf{x} around $\widehat{\mathbf{b}}$.²⁸ Moreover, assume that costs are increasing as we scale the distance from $\widehat{\mathbf{b}}$:

Assumption 4. If $\mathbf{x} \in \mathbb{R}^n$ while $s > 1$ then

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

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

$$\begin{aligned} \max_{\mathbf{b}} F(\mathbf{a}^\top \mathbf{a}) & \quad (\text{IT-G}) \\ \text{s.t. } K(\mathbf{b}; \widehat{\mathbf{b}}) & \leq C. \end{aligned}$$

As in Section 4, define

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

Our main result on this is:

Proposition 5. Assume \mathbf{M} is invertible and Assumptions 3-4 hold. Suppose \mathbf{b}^* solves (IT-G). Then for generic $\widehat{\mathbf{b}}$, we have that \underline{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 is an increasing function of the Euclidean distance of the new \mathbf{b} from the base point $\widehat{\mathbf{b}}$; Assumptions 3-4 guarantee this property. Second, the proof of Proposition 1 can easily be modified to dispense with normality of \mathbf{G} . The only cost is that the result must be stated in terms of singular values of \mathbf{M} rather than eigenvalues of \mathbf{G} . The disadvantage of that is that the singular

²⁷The only difference with the case in which $\mathbf{M} = \mathbf{I} - \beta\mathbf{G}$ and \mathbf{G} is normal is that the singular values s_l of \mathbf{M} will not be equal to $1 - \beta\lambda_l$, where λ_l are eigenvalues of \mathbf{G} , and the singular vectors of \mathbf{M} will not be eigenvectors of \mathbf{G}

²⁸The quadratic cost function we consider in the basic setting satisfies Assumption 3

values of \mathbf{M} depend on β , so we can no longer make clean statements about how strategic complements and substitutes differ. The advantage is that there is still a clean summary, of the priorities of the planner in terms of classical invariants of a matrix capturing the strategic interactions.

6.3. Nonlinear systems. In this section we will study the volatility-control problem of Section 5.1 for a nonlinear system. Consider an economic system in which the endogenous investment profile \mathbf{a} solves the following system:

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

where f_i is a differentiable function. Fix a solution $\hat{\mathbf{a}}$ of this system corresponding to $\hat{\mathbf{b}}$, 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 $\hat{\mathbf{b}}$, so that $\hat{\mathbf{b}}$ changes to $\hat{\mathbf{b}} + \check{\mathbf{x}}$. 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

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

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

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

Define $\mathbf{G} = \mathbf{F}\mathbf{H}$. Then

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

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

$$\check{\mathbf{a}} = \mathbf{S}^{-1}\mathbf{F}\check{\mathbf{x}} + O(\|\check{\mathbf{x}}\|^2). \tag{10}$$

We obtain two insights. First, in this non-linear economy, a shock in the idiosyncratic characteristics of the l^{th} principal component of \mathbf{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,

$$\text{Var} \left(\sum_i a_i \right) = \mathbb{E} [\check{\mathbf{a}}^\top \check{\mathbf{a}}] = \mathbb{E} [\check{\mathbf{a}}^\top \check{\mathbf{a}}] = \mathbb{E} [(\mathbf{F}\check{\mathbf{x}})^\top \mathbf{S}^{-2}(\mathbf{F}\check{\mathbf{x}})] + O(\mathbb{E}\|\check{\mathbf{x}}\|^4)$$

In the last line we plugged in (10). We have taken the terms that appear such as $\mathbb{E}[\check{\mathbf{a}}^\top \mathbf{v}]$, where $\mathbf{v} = O(\|\check{\mathbf{x}}\|_2)$, and bounded them by $O(\mathbb{E}\|\check{\mathbf{x}}\|^4)$. This follows because $\mathbb{E}[\check{\mathbf{a}}]$ is close to 0, with an error of order $O(\|\mathbf{x}\|^2)$.

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

7. CONCLUDING REMARKS

We have developed a framework to study optimal interventions when individuals interact strategically with their neighbors. We solve this class of intervention problems by exploiting singular value decompositions of matrices capturing 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 in our analysis have some connections to ones that have been studied in the network literature but, more importantly, offer a new set of network statistics that are simple to describe and may be of interest to empirical researchers. For instance, in the basic peer effects application, the optimal targets focus on the first principal component: this corresponds to eigenvector centrality of the matrix of interactions, a widely-studied network statistic. But it is the last principal component—i.e., the one corresponding to the least singular value of the network—that matters most under strategic substitutes. The last singular component captures the “local” structure of the network, explaining how best to avoid crowd-out among neighbors. This aspect of the network, to our knowledge, has not been identified as important in network games, and yet we show it is the essential one in games of strategic substitutes.

In general other principal components matter, and they capture a range of structural aspects of the network, ranging from more global summary statistics (for large eigenvalues of the matrix of interactions) to more local (for smaller ones). The fact that principal components have been a powerful and illuminating tool in applied mathematics and across many fields of economics 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.

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.
- 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*.
- BELHAJ, M. AND F. DEROIAN (2017): “Contracting on networks,” 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.
- 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.
- CANDOGAN, O., K. BIMPIKIS, AND A. OZDAGLAR (2012): “Optimal Pricing in Networks with Externalities,” *Operations Research*, 60, 883–905.
- 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.

- 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.
- 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. AND S. GOYAL (2009): “Influencing the Infleuncers: 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 S. MORRIS (2017): “Higher-Order Expectations,” Available at SSRN: <http://ssrn.com/abstract=2979089>.
- GOLUB, G. H. AND C. F. VAN LOAN (1996): *Matrix Computations*, The Johns Hopkins University Press.
- 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.

- 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*.
- 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., ed. (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.
- VALENTE, T. (2012): “Network Interventions,” *Science*, 337, 49–53.
- VIVES, X. (2001): *Oligopoly Pricing*, Cambridge, MA: The MIT 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 orthonormal transformation \mathbf{V}^T which takes variables to their “underlined” coordinates:

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

and

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

By defining

$$\alpha_l = \frac{1}{2(1 - \beta \lambda_l(\mathbf{G}))^2}, \tag{11}$$

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

$$\max_{\underline{\mathbf{b}}} \sum_{l=1}^n \alpha_l \underline{b}_l^2 \tag{IT-SVD}$$

$$\text{s.t. } \sum_{l=1}^n [b_l - \widehat{b}_l]^2 \leq C.$$

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

$$\begin{aligned} \max_{\underline{b}} \quad & \sum_{l=1}^n \alpha_l \widehat{b}_l^2 [b_l + 1]^2 & (\text{IT-SVD-REL}) \\ \text{s.t.} \quad & \sum_{l=1}^n \widehat{b}_l^2 b_l^2 \leq C. \end{aligned}$$

Note that, for all l the α_l defined by (11) are well-defined (by Assumption 1) and strictly positive. This has two implications.

First, if \underline{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 b_l^2 . Either the increase or the decrease is guaranteed to increase the corresponding $[b_l + 1]^2$ (since the α_l are all strictly positive).

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

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

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

$$0 = \frac{\partial \mathcal{L}}{\partial b_l} = 2\widehat{b}_l^2 [\alpha_l(1 + b_l^*) - \mu b_l^*] \quad l = 1, 2, \dots, n. \quad (12)$$

We will take a generic $\widehat{\underline{b}}$ such that $\widehat{b}_l \neq 0$ for each l . If for some l we had $\mu = \alpha_l$ then the right-hand side of (12) would be $2\widehat{b}_l^2 \alpha_l$, which, by the generic assumption we just made and the positivity of α_l , would contradict (12). Thus the following holds with a nonzero denominator:

$$b_l^* = \frac{\alpha_l}{\mu - \alpha_l}. \quad (13)$$

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

A.1.2. *Proof of Proposition 2:* Consider first the case $\beta > 0$. First note that

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

where the equality follows by (13) from the proof of Proposition 1, and the inequality follows because $\mu > \alpha_1$ (a fact also argued in the proof of Proposition 1). Hence, because $\alpha_l \leq \alpha_2$ for any $l \neq 1$, for $\underline{b}_l/\underline{b}_1 < \epsilon$ it is sufficient that

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

Now we will fix δ and argue that if C exceeds some level, then $1 - \frac{\alpha_1}{\mu} < \delta$. For convenience, we write the condition on δ as $\frac{\mu}{\alpha_1} - 1 < \delta \frac{\mu}{\alpha_1}$. At the optimum, the constraint in (IT-SVD-REL) binds (as argued in the proof of Proposition 1), i.e.,

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

Then note that if $\frac{\mu}{\alpha_1} - 1 \geq \delta \frac{\mu}{\alpha_1}$ then

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

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

Taking the contrapositive, if $C > \frac{\|\widehat{\mathbf{b}}\|_2^2}{\delta^2}$ then $\frac{\mu}{\alpha_1} - 1 < \delta \frac{\mu}{\alpha_1}$. Using the definition of δ in terms of ϵ completes the first part of the statement.

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 $C \rightarrow \infty$, we have $\underline{b}_l^*/\underline{b}_1^* \rightarrow 0$ for all $l \neq 1$. By definition of $\widehat{\mathbf{b}}$, this translates into

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

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

$$\frac{[b_l^* - \widehat{b}_l]^2}{[b_1^* - \widehat{b}_1]^2} \rightarrow 0.$$

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

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

Now, dividing the previous equation by $[b_1^* - \widehat{b}_1]^2$, and using the previous statement about the limit, makes it clear that $[b_l^* - \widehat{b}_l]^2 / C \rightarrow 1$, and thus, from the constraint, $[b_l^* - \widehat{b}_l]^2 / C \rightarrow 0$ for $l \neq 1$.

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

A.2. Pricing game.

Lemma 1. Consider the pricing game and consider a change from $\widehat{\Omega}$ to Ω the change in the expected consumer surplus, producer surplus and welfare are given by expressions 7-9.

Proof. Consumer surplus equals the equilibrium utility of the representative consumers, i.e., $CS(\mathbf{p}) = \frac{1}{4} \sum_f [Q_f(\mathbf{p})]^2$, and therefore

$$\begin{aligned} \mathbb{E}_{\Omega} [CS(\mathbf{p})] &= \frac{1}{4} \mathbb{E}_{\Omega} [[\gamma \mathbf{1}_F - \mathbf{T}^T \mathbf{p}]^T [\gamma \mathbf{1}_F - \mathbf{T}^T \mathbf{p}]] \\ &= \mathbb{E}_{\Omega} [\gamma^2 \mathbf{1}_F^T \mathbf{1} - 2\gamma \mathbf{1}_F^T \mathbf{T}^T \mathbf{p}] + \mathbb{E}_{\Omega} [\mathbf{p}^T \Lambda \mathbf{p}], \end{aligned}$$

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

$$\begin{aligned} \mathbb{E}_{\Omega} [CS(\mathbf{p})] - \mathbb{E}_{\widehat{\Omega}} [CS(\mathbf{p})] &= \mathbb{E}_{\Omega} [\mathbf{p}^T \Lambda \mathbf{p}] - \mathbb{E}_{\widehat{\Omega}} [\mathbf{p}^T \Lambda \mathbf{p}] \\ &= [\mathbf{I} + \Lambda]^{-2} \Lambda [\mathbb{E}_{\Omega} [\mathbf{b}^T \mathbf{b}] - \mathbb{E}_{\widehat{\Omega}} [\mathbf{b}^T \mathbf{b}]] \\ &= [\mathbf{I} + \Lambda]^{-2} \Lambda [\mathbb{E}_{\Omega} [\mathbf{c}^T \mathbf{c}] - \mathbb{E}_{\widehat{\Omega}} [\mathbf{c}^T \mathbf{c}]], \end{aligned}$$

which is equivalent to (7). Next, $PS(\mathbf{p}) = \sum_i \Pi_i(\mathbf{p})$ and so

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

Hence

$$\begin{aligned}
\mathbb{E}_{\Omega} [PS(\mathbf{p})] - \mathbb{E}_{\hat{\Omega}} [PS(\mathbf{p})] &= \mathbb{E}_{\Omega} [\underline{\mathbf{p}}^{\top} \underline{\mathbf{p}} + \underline{\mathbf{c}}^{\top} \underline{\mathbf{c}} - 2\underline{\mathbf{p}}^{\top} \underline{\mathbf{c}}] \\
&= [(\mathbf{I} + \mathbf{\Lambda})^{-2} + \mathbf{I} - 2(\mathbf{I} + \mathbf{\Lambda})^{-1}] [\mathbb{E}_{\Omega}[\underline{\mathbf{c}}^{\top} \underline{\mathbf{c}}] - \mathbb{E}_{\hat{\Omega}}[\underline{\mathbf{c}}^{\top} \underline{\mathbf{c}}]] \\
&= \frac{\mathbf{\Lambda}^2}{(\mathbf{I} + \mathbf{\Lambda})^{-2}} [\mathbb{E}_{\Omega}[\underline{\mathbf{c}}^{\top} \underline{\mathbf{c}}] - \mathbb{E}_{\hat{\Omega}}[\underline{\mathbf{c}}^{\top} \underline{\mathbf{c}}]],
\end{aligned}$$

which is equivalent to (8). The expression for (9) follows by combining (7) and (8). \square

A.3. 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 $\|\mathbf{b} - \hat{\mathbf{b}}\|_2$. Indeed, for any \mathbf{x} , we can find an orthonormal transformation \mathbf{O} that maps \mathbf{x} to $\|\mathbf{x}\|_2 \mathbf{1}$. Thus, by Assumption 3, K is fully determined by its values on $K(\hat{\mathbf{b}} + s\mathbf{1}; \hat{\mathbf{b}})$ as s ranges over $[0, \infty]$. We know by Assumption 4 that these are increasing in s , let's say according to some function $k : \mathbb{R} \rightarrow \mathbb{R}$. Putting these facts together we can see that $K(\hat{\mathbf{b}} + \mathbf{x}; \hat{\mathbf{b}}) = k(\|\mathbf{x}\|_2)$.

Thus (IT-G) is equivalent to

$$\begin{aligned}
&\max_{\mathbf{b}} \mathbf{a}^{\top} \mathbf{a} && \text{(IT-G-S)} \\
&\text{s.t. } \|\mathbf{b} - \hat{\mathbf{b}}\|_2^2 \leq C,
\end{aligned}$$

which is identical to the problem we studied in section 4 except for the lack of normality. The proof of proposition 5 is at this stage analogous to the proof of Proposition 1. The only difference is that we define $\underline{\mathbf{a}} = \mathbf{U}^{\top} \mathbf{a}$ and $\underline{\mathbf{b}} = \mathbf{V}^{\top} \mathbf{b}$, and define α_l in the proof of that result to be s_l , the corresponding singular value, avoiding the eigenvalues altogether.