

TARGETING INTERVENTIONS IN NETWORKS

ANDREA GALEOTTI, BENJAMIN GOLUB, AND SANJEEV GOYAL

ABSTRACT. We study games in which a network mediates strategic spillovers and externalities among the players. How does a planner optimally target interventions that change individuals' private returns to investment? We analyze this question by decomposing any intervention into orthogonal *principal components*, which are determined by the network and are ordered according to their associated eigenvalues. There is a close connection between the nature of spillovers and the representation of various principal components in the optimal intervention. In games of strategic complements (substitutes), interventions place more weight on the top (bottom) principal components, which reflect more global (local) network structure. For large budgets, optimal interventions are simple – they involve a single principal component.

1. INTRODUCTION

We study games among agents embedded in a network. The action of each agent – e.g., a level of investment or effort – directly affects a subset of others, called *neighbors* of that agent. This happens through two channels: spillover effects on others' incentives, as well as non-strategic externalities. A utilitarian planner with limited resources can intervene to change individuals' incentives for taking the action. Our goal is to understand how the planner can best target such interventions in view of the network and other primitives of the environment.

We now lay out the elements of the model in more detail. Individuals play a simultaneous-move game with continuous actions. An agent's action confers *standalone* benefits on that agent independent of anyone else's action, but it also creates spillovers. The intensity of these spillovers is described by a network, with the strength of a link between two individuals reflecting how strongly the action of one affects the marginal benefits experienced by the other. The effects may take the form of strategic complements or strategic substitutes. In

We are grateful to the co-editor, Dirk Bergemann, and five anonymous referees for helpful comments. We have also benefited from conversations with Francis Bloch, Drew Fudenberg, Eric Maskin, Matthew Jackson, Asuman Ozdaglar, Francesca Parise, Omer Tamuz, Eduard Talamàs, John Urschel, Xavier Vives, Rakesh Vohra, Alex Wolitzky, Leeat Yariv. Ria Granzier-Nakajima, Joerg Kalbfuss, Gustavo Paez, and Eduard Talamàs provided exceptional research assistance. We thank Sihua Ding, Joerg Kalbfuss, Fakhteh Saadatniaki, Alan Walsh, and Yves Zenou for detailed comments on earlier drafts. Andrea Galeotti gratefully acknowledges financial support from European Research Council through the ERC-consolidator grant (award no. 724356) and the European University Institute through the Internal Research Grant. Benjamin Golub gratefully acknowledges financial support from The Pershing Square Fund for Research on the Foundations of Human Behavior and the National Science Foundation (SES-1658940, SES-1629446). Galeotti: Department of Economics, London Business School, agaleotti@london.edu. 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.

addition to standalone benefits and incentive spillovers, there may be positive or negative externalities imposed by network neighbors on each other.¹ Before this game is played, the planner can target some individuals and alter their standalone marginal benefits from status quo levels. The cost of the intervention is increasing in the magnitude of the change and is separable across individuals. The planner seeks to maximize the utilitarian welfare under equilibrium play of the game, subject to a budget constraint. Our results characterize the optimal intervention policy, showing how it depends on the network, the nature of spillovers, the status quo incentives, and the budget.

An intervention on one individual has direct and indirect effects on the incentives of others. These effects depend on the network and on whether the game features strategic substitutes or complements. For example, suppose the planner increases a given individual's standalone marginal benefits to effort, thereby increasing effort by the targeted individual. If actions are strategic complements, this will push up the incentives of the targeted individual's neighbors. That will increase the efforts of the neighbors of these neighbors, and so forth, creating aligned feedback effects throughout the network. In contrast, under strategic substitutes, the same intervention will discourage the individual's neighbors from exerting effort. However, the effect on those neighbors' neighbors will be positive – i.e., in the same direction as the effect on the targeted agent. This interplay between spillovers and network structure makes targeting interventions a complex problem.

At the heart of our approach is a particular way to organize the spillover effects in terms of the *principal components*, or eigenvectors, of the matrix of interactions. Any change in the vector of standalone marginal returns can be expressed in a basis of these principal components. This basis has three special properties: (a) when standalone marginal returns are exogenously changed in the direction of a principal component, the effect is to change equilibrium actions *in the same direction*; (b) the magnitude of the effect is a multiple of the magnitude of the exogenous change, and the multiplier is determined by an eigenvalue of the network corresponding to that principal component; (c) the principal components are orthogonal, so that the effects along various principal components can be treated separately. The three properties we have listed permit us to express the effect of interventions on actions, and on welfare, in a way that facilitates a simple characterization of optimal interventions.

Our main result, Theorem 1, characterizes the optimal intervention in terms of how similar it is to various principal components – or, in other words, how strongly represented various

¹This framework encompasses a number of well-known economic examples from the literature: spillovers in educational/criminal effort (Ballester, Calvó-Armengol, and Zenou, 2006), research collaboration among firms (Goyal and Moraga-Gonzalez, 2001), local public goods (Bramouille and Kranton, 2007), investment games and beauty contests (Angeletos and Pavan, 2007; Morris and Shin, 2002), and peer effects in smoking (Jackson et al., 2017).

principal components are in it.² Building on this characterization, Corollary 1 describes how the nature of the strategic interaction shapes which principal components figure most prominently in the optimal intervention. The principal components can be ordered by their associated eigenvalues (from high to low). In games of strategic complements, the optimal intervention is most similar to the first principal component – the familiar *eigenvector centrality* – and progressively less similar as we move down the principal components. In games of strategic substitutes, the order is reversed: the optimal intervention is most similar to the *last* principal component. The “higher” principal components capture the more global structure of the network: this is important for taking advantage of the aligned feedback effects arising under strategic complementarities. The “lower” principal components capture the local structure of the network: they help the planner to target the intervention so that it does not cause crowding out between adjacent neighbors: this is an important concern when actions are strategic substitutes.

We then turn to the study of *simple* optimal interventions, i.e., ones where the relative intervention (change in the standalone marginal benefit) for each node is determined by a single network statistic of that node, and invariant to other primitives such as status quo incentives. Propositions 1 and 2 show that for large enough budgets the optimal intervention is simple: in games of strategic complements, the optimal intervention vector is proportional to the first principal component, while in games of strategic substitutes, it is proportional to the last one.³ Moreover, the network structure determines how large the budget must be for optimal interventions to be simple. In games of strategic complements (substitutes), the important statistic is the gap between the top (bottom) two eigenvalues. When this gap is large, even at moderate budgets the intervention is simple.

Theorem 1 obtains in a setting where the planner knows the status quo standalone marginal benefits of all individuals. Our methods can also be used to study optimal interventions in a setting where the planner does not know these benefits but knows only their distribution. In such a setting, the impact of an intervention on expected social welfare is determined by how it alters the first and second moments of the benefits distribution. Propositions 3 and 4 characterize optimal interventions and show that the key insights about the order of targeting of principal components extend.

We now place the paper in the context of the literature. The intervention problem we study concerns optimal policy in the presence of externalities. Research over the past two decades has deepened our understanding of the empirical structure of networks and the

²We use the standard notion of cosine similarity: the similarity of two vectors is the cosine of the angle between them in the plane they jointly define.

³In similarity terms, this means that the optimal intervention has a cosine similarity of nearly 1 to the first or last principal component (depending on the case), and a similarity of nearly 0 to all other principal components.

theory of how networks affect strategic behavior.⁴ This has led to the study of how policy design should incorporate information about networks. Network interventions are currently an active subject of research not only in economics but also in related disciplines such as computer science, sociology and public health.⁵ The main contribution of this paper is methodological. It lies in (i) using the principal components approach to decompose the effect of an intervention on social welfare and (ii) using the structure afforded by this decomposition to characterize optimal interventions. Of special interest is the close relation between the strategic structure (complements or substitutes) and the appropriate principal components to target.⁶

The rest of the paper is organized as follows. Section 2 presents the optimal intervention problem. Section 3 sets out how we apply a principal component decomposition to our game. Section 4 characterizes optimal interventions. Section 5 studies a setting where the planner has incomplete information about agents' standalone marginal returns. Section 6 concludes. Appendix A contains the proofs of the main results. The Online Appendix presents the proofs of Propositions 3 and 4 and discusses a number of extensions.

2. THE MODEL

We consider a simultaneous-move game among individuals $\mathcal{N} = \{1, \dots, n\}$, where $n \geq 2$. Individual i chooses an action, $a_i \in \mathbb{R}$. The vector of actions is denoted by $\mathbf{a} \in \mathbb{R}^n$. The payoff to individual i depends on this vector, \mathbf{a} , the *network* with adjacency matrix \mathbf{G} , and other parameters, described below:

$$U_i(\mathbf{a}, \mathbf{G}) = \underbrace{a_i \left(b_i + \beta \sum_{j \in \mathcal{N}} g_{ij} a_j \right)}_{\text{returns from own action}} - \underbrace{\frac{1}{2} a_i^2}_{\text{private costs of own action}} + \underbrace{P_i(\mathbf{a}_{-i}, \mathbf{G}, \mathbf{b})}_{\text{pure externalities}}. \quad (1)$$

The private marginal returns from increasing the action a_i depend both on i 's own action, a_i , and on others' actions. The coefficient $b_i \in \mathbb{R}$ corresponds to the part of i 's marginal return that is independent of others' actions, and is thus called i 's *standalone marginal return*. The contribution of others' actions to i 's marginal return is given by the term $\beta \sum_{j \in \mathcal{N}} g_{ij} a_j$. Here $g_{ij} \geq 0$ is a measure of the strength of the interaction between i and j ; we assume that for every $i \in \mathcal{N}$, $g_{ii} = 0$ – there are no self-loops in the network \mathbf{G} . The parameter β captures

⁴See, for example, Goyal, Moraga, and van der Leij (2006), Ballester, Calvó-Armengol, and Zenou (2006), Bramoullé, Kranton, and d'Amours (2014), and Galeotti, Goyal, Jackson, Vega-Redondo, and Yariv (2010)

⁵For a general introduction to the subject, see Rogers (1983), Kempe, Kleinberg, and Tardos (2003), Borgatti (2006), and Valente (2012). Within economics, a prominent early contribution is Ballester, Calvó-Armengol, and Zenou (2006); recent contributions include Banerjee, Chandrasekhar, Duflo, and Jackson (2013), Belhaj and Deroian (2017), Bloch and Querou (2013), Candogan, Bimpikis, and Ozdaglar (2012), Demange (2017), Fainmesser and Galeotti (2017), Galeotti and Goyal (2009), Galeotti and Rogers (2013), Leduc, Jackson, and Johari (2017), and Akbarpour, Malladi, and Saberi (2017).

⁶Online Appendix Section OA2.1 presents a discussion of the relationship between principal components and other network measures that have been studied in the literature.

strategic interdependencies. If $\beta > 0$, then actions are strategic complements; if $\beta < 0$, then actions are strategic substitutes. The function $P_i(\mathbf{a}_{-i}, \mathbf{G}, \mathbf{b})$ captures *pure externalities* – that is, spillovers that do not affect the best response. The first-order condition for individual i 's action to be a best response is:

$$a_i = b_i + \beta \sum_{j \in \mathcal{N}} g_{ij} a_j.$$

Any Nash equilibrium action profile \mathbf{a}^* of the game satisfies

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

We now make two assumptions about the network and the strength of strategic spillovers. Recall that the spectral radius of a matrix is the maximum of its eigenvalues' absolute values.

Assumption 1. The adjacency matrix \mathbf{G} is symmetric.⁷

Assumption 2. The spectral radius of $\beta \mathbf{G}$ is less than 1,⁸ and all eigenvalues of \mathbf{G} are distinct (the latter condition holds generically).

Assumption 2 ensures that (2) is a necessary and sufficient condition for a solution, and also ensures the uniqueness and stability of the Nash equilibrium.⁹ Under these assumptions, the unique Nash equilibrium of the game can be characterized by

$$\mathbf{a}^* = [\mathbf{I} - \beta \mathbf{G}]^{-1} \mathbf{b}. \quad (3)$$

The utilitarian social welfare at equilibrium is given by the sum of the equilibrium utilities:

$$W(\mathbf{b}, \mathbf{G}) = \sum_{i \in \mathcal{N}} U_i(\mathbf{a}^*, \mathbf{G}).$$

The planner aims to maximize the utilitarian social welfare at equilibrium. She does this by changing the status quo standalone marginal returns $\hat{\mathbf{b}}$, to new values, \mathbf{b} , subject to a budget constraint on the cost of her intervention. The timing is as follows. The planner moves first and chooses her intervention, and then individuals simultaneously choose actions. The incentive-targeting (IT) problem is given by

$$\begin{aligned} & \max_{\mathbf{b}} W(\mathbf{b}, \mathbf{G}) && \text{(IT)} \\ & \text{s.t.: } \mathbf{a}^* = [\mathbf{I} - \beta \mathbf{G}]^{-1} \mathbf{b}, \\ & K(\mathbf{b}, \hat{\mathbf{b}}) = \sum_{i \in \mathcal{N}} (b_i - \hat{b}_i)^2 \leq C, \end{aligned}$$

⁷We extend our analysis to more general \mathbf{G} in the Online Appendix Section OA3.2.

⁸An equivalent condition is for $|\beta|$ to be less than the reciprocal of the spectral radius of \mathbf{G} .

⁹See Ballester et al. (2006) and Bramoullé et al. (2014) for detailed discussions of this assumption and the interpretation of the solution given by (3).

where C is a given budget. The function K is an adjustment cost of implementing an intervention. The marginal costs of altering the b_i are separable across individuals and increasing in the magnitude of the change for each individual. For discussions of the adjustment costs and extensions on more general planner cost functions, see the Online Appendix Section OA3.3. In the Online Appendix Section OA3.4, we study a setting in which a planner provides monetary payments to individuals that induce them to change their actions. We show that the resulting optimal intervention problem has the same mathematical structure as the one we study in our basic model.

We present two economic applications to illustrate the scope of our model. The first example is a classical investment game, and the second example is a game of providing a local public good.

Example 1 (The investment game). Individual i makes an investment a_i at a cost $\frac{1}{2}a_i^2$. The private marginal return on that investment is $b_i + \beta \sum_{j \in \mathcal{N}} g_{ij}a_j$, where b_i is individual i 's standalone marginal return and $\sum_{j \in \mathcal{N}} g_{ij}a_j$ is the aggregate local effort. The utility of i is

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

The case with $\beta > 0$ reflects investment complementarities, as in Ballester et al. (2006). Here, an individual's marginal returns are enhanced when his neighbors work harder; this creates both strategic complementarities and positive externalities. The case of $\beta < 0$ corresponds to strategic substitutes and negative externalities; this can be microfounded via a model of competition in a market after the investment decisions a_i have been made, as in Goyal and Moraga-Gonzalez (2001). A manager who observes the network of strategic interactions—for instance, which agents work together on joint projects—can intervene by changing levels of monitoring or encouragement relative to a status quo level.

It can be verified that the equilibrium utilities, $U_i(\mathbf{a}^*, \mathbf{G})$, and the utilitarian social welfare at equilibrium, $W(\mathbf{b}, \mathbf{G})$, are as follows:

$$U_i(\mathbf{a}^*, \mathbf{G}) = \frac{1}{2}(a_i^*)^2 \text{ and } W(\mathbf{b}, \mathbf{G}) = \frac{1}{2}(\mathbf{a}^*)^\top \mathbf{a}^*.$$

Example 2 (Local public goods). We next consider a local public goods problem in a framework that follows the work of Bramoulle and Kranton (2007), Galeotti and Goyal (2010), and Allouch (2015, 2017). In a local public goods problem, each agent makes a costly contribution, which brings her closer to an ideal level of public goods but also raises the levels enjoyed by her neighbors. Examples include (i) contributions to improve physical neighborhoods, such as residents clearing snow or workers keeping areas on a factory floor clean and safe and (ii) knowledge workers acquiring non-rival information (e.g., about job applicants) that can be shared with colleagues. In example (i), the network governing

spillovers is given by physical proximity, while in example (ii) it is given by organizational overlap. To be concrete, in example (i), agents receive some level of municipal services at the status quo. They augment it with their own effort, and benefit (with a discount) from the efforts contributed by neighbors. A city councilor who observes the network structure of physical proximity among houses can intervene to change the status quo allocation of services, tailoring it to improve incentives.

Formally, suppose that if each i contributes effort a_i to the public good, then the amount of public good i experiences is

$$x_i = \tilde{b}_i + a_i + \tilde{\beta} \sum_{j \in \mathcal{N}} g_{ij} a_j,$$

where $0 < \tilde{\beta} < 1$. The utility of i is

$$U_i(\mathbf{a}, \mathbf{G}) = -\frac{1}{2}(\tau - x_i)^2 - \frac{1}{2}a_i^2,$$

where $\tilde{b}_i < \tau$.

We now interpret these formulas. The optimal level of public good in the absence of any costs is τ ; this can be thought of as the maximum that can be provided. Individual i has access to a base level \tilde{b}_i of the public good. Each agent can expend costly effort, a_i , to augment this base level, and receive $\tilde{b}_i + a_i$. If i 's neighbor j expends effort, a_j , then i has access to an additional $\tilde{\beta} g_{ij} a_j$ units of the public good, where $\tilde{\beta} < 1$.

This is a game of strategic substitutes and positive externalities. Performing the change of variables $b_i = [\tau - \tilde{b}_i]/2$ and $\beta = -\tilde{\beta}/2$ (with the status quo equal to $\hat{b}_i = [\tau - \tilde{b}_i]/2$) yields a best-response structure exactly as in condition (2). The aggregate equilibrium utility is $W(\mathbf{b}, \mathbf{G}) = -(\mathbf{a}^*)^\top \mathbf{a}^*$.

These two canonical examples share a technically convenient property:

Property A. The aggregate equilibrium utility is proportional to the sum of the squares of the equilibrium actions, that is, $W(\mathbf{b}, \mathbf{G}) = w \cdot (\mathbf{a}^*)^\top \mathbf{a}^*$ for some $w \in \mathbb{R}$, where \mathbf{a}^* is the Nash equilibrium of the network game.

Online Appendix Section OA2.2 discusses a network beauty contest game inspired by Morris and Shin (2002) and Angeletos and Pavan (2007) which also satisfies this property. While Property A facilitates analysis, it is not essential. Online Appendix Section OA3.1 extends the analysis to cover important cases where this property does not hold.

3. PRINCIPAL COMPONENTS

This section introduces a basis for the space of standalone marginal returns and actions in which, under our assumptions on \mathbf{G} , strategic effects and the welfare function of interest to the planner both take a simple form.

Fact 1. If \mathbf{G} satisfies Assumption 1, then $\mathbf{G} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$, where:

1. $\mathbf{\Lambda}$ is an $n \times n$ diagonal matrix whose diagonal entries $\Lambda_{\ell\ell} = \lambda_\ell$ are the eigenvalues of \mathbf{G} (which are real numbers), ordered from greatest to least: $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$.
2. \mathbf{U} is an orthogonal matrix. The ℓ^{th} column of \mathbf{U} , which we call \mathbf{u}^ℓ , is a real eigenvector of \mathbf{G} , namely the eigenvector associated to the eigenvalue λ_ℓ , which is normalized in the Euclidean norm: $\|\mathbf{u}^\ell\| = 1$.

For generic \mathbf{G} , the decomposition is uniquely determined, except that any column of \mathbf{U} is determined only up to multiplication by -1 .

An important interpretation of this diagonalization is as a decomposition into *principal components*. First, consider the symmetric rank-one matrix that best approximates \mathbf{G} in the squared-error sense—in other words, the vector \mathbf{u} such that

$$\sum_{i,j \in \mathcal{N}} (g_{ij} - u_i u_j)^2$$

is minimized. The minimizer turns out to be a scaling of the eigenvector \mathbf{u}^1 . Now, if we consider the “residual” matrix $\mathbf{G}^{(2)} = \mathbf{G} - \mathbf{u}^1(\mathbf{u}^1)^\top$, we can perform the same rank-one approximation and obtain the second eigenvector \mathbf{u}^2 . Proceeding further in this way gives us a sequence of vectors that constitute an orthonormal basis. At each step, the next vector generates the rank-one matrix that best summarizes the remaining structure in the matrix \mathbf{G} .¹⁰

Figure 1 illustrates some eigenvectors/principal components of a circle network with 14 nodes and with links all having equal weight given by 1. For each principal component, the color of a node indicates the sign of the entry of that node in that principal component (the color red means negative), while the size of a node indicates the absolute value of that entry.¹¹ A general feature that is worth noting is that the entries of the top principal components (with smaller values of ℓ) are similar among neighboring nodes, while the bottom principal components (with larger values of ℓ) tend to be negatively correlated among neighboring nodes.¹²

3.1. Analysis of the game using principal components. For any vector $\mathbf{z} \in \mathbb{R}^n$, let $\underline{z} = \mathbf{U}^\top \mathbf{z}$. We will refer to \underline{z}_ℓ as the projection of \mathbf{z} onto the ℓ^{th} principal component, or the magnitude of \mathbf{z} in that component. Substituting the expression $\mathbf{G} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$ into equation (2), which characterizes equilibrium, we obtain

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

¹⁰See Spielman (2007), especially subsection 16.5.1, on this interpretation. For book-length treatments of spectral graph theory, see Cvetkovic et al. (1997) and Chung and Graham (1997).

¹¹The circle network is invariant to rotations (cyclic permutations) of the nodes and so the eigenvectors are determined only up to a rotation.

¹²For a formal treatment of this, see Davies et al. (2001) and Urschel (2018).

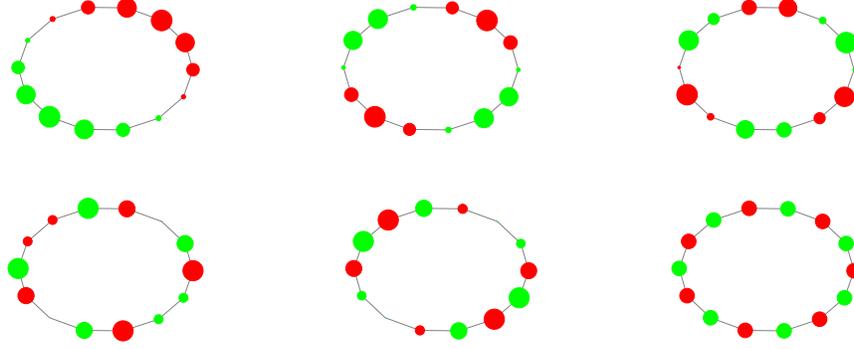


FIGURE 1. (Top) Eigenvectors 2, 4, 6. (Bottom) Eigenvectors 10, 12, 14.

Multiplying both sides of this equation by \mathbf{U}^\top gives us an analogue of (3) characterizing the solution of the game:

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

This system is diagonal, and the ℓ^{th} diagonal entry of $[\mathbf{I} - \beta\mathbf{\Lambda}]^{-1}$ is $\frac{1}{1 - \beta\lambda_\ell}$. Hence, for every $\ell \in \{1, 2, \dots, n\}$,

$$\underline{a}_\ell^* = \frac{1}{1 - \beta\lambda_\ell} \underline{b}_\ell. \quad (4)$$

The principal components of \mathbf{G} constitute a basis in which strategic effects are easily described. The equilibrium action \underline{a}_ℓ^* in the ℓ^{th} principal component of \mathbf{G} is the product of an amplification factor (determined by the strategic parameter β and the eigenvalue λ_ℓ) and \underline{b}_ℓ , which is simply the projection of \mathbf{b} onto that principal component. Under Assumption 2, for all ℓ we have $1 - \beta\lambda_\ell > 0$.¹³ Moreover, when $\beta > 0$ ($\beta < 0$), the amplification factor is decreasing (increasing) in ℓ .

We can also use this to give a formula for equilibrium actions in the original coordinates:

$$a_i^* = \sum_{\ell=1}^n \frac{1}{1 - \beta\lambda_\ell} u_i^\ell \underline{b}_\ell.$$

We close with a definition that will allow us to describe optimal interventions in terms of a standard measure of their similarity to various principal components.

Definition 1. The cosine similarity of two nonzero vectors \mathbf{z} and \mathbf{y} is

$$\rho(\mathbf{z}, \mathbf{y}) = \frac{\mathbf{z} \cdot \mathbf{y}}{\|\mathbf{z}\| \|\mathbf{y}\|}.$$

This is the cosine of the angle between the two vectors in the plane determined by \mathbf{y} and \mathbf{z} . When $\rho(\mathbf{z}, \mathbf{y}) = 1$, vector \mathbf{z} is a positive scaling of \mathbf{y} . When $\rho(\mathbf{z}, \mathbf{y}) = 0$, vectors \mathbf{z} and \mathbf{y} are orthogonal. When $\rho(\mathbf{z}, \mathbf{y}) = -1$, vector \mathbf{z} is a negative scaling of \mathbf{y} .

¹³Assumption 2 on the spectral radius implies that $\beta\mathbf{\Lambda}$ has no entries larger than 1.

4. OPTIMAL INTERVENTIONS

This section develops a characterization of optimal interventions and studies their properties.

We begin by dispensing with a straightforward case of the planner's problem. Recall that under Property A, the planner's payoff as a function of the equilibrium actions \mathbf{a}^* is $W(\mathbf{b}, \mathbf{G}) = w \cdot (\mathbf{a}^*)^\top \mathbf{a}^*$. If $w < 0$, the planner wishes to minimize the sum of the squares of the equilibrium actions. In this case, when the budget is large enough, that is, $C \geq \|\hat{\mathbf{b}}\|^2$, the planner can allocate resources to ensure that individuals have a zero target action by setting $b_i = 0$ for all i . It follows from the best-response equations that all individuals choose action 0 in equilibrium, and so the planner achieves the first-best.¹⁴ The next assumption implies that the planner's bliss point cannot be achieved, so that there is an interesting optimization problem:

Assumption 3. Either $w < 0$ and $C < \|\hat{\mathbf{b}}\|$, or $w > 0$.

Let \mathbf{b}^* solve the incentive-targeting problem (IT), and let $\mathbf{y}^* = \mathbf{b}^* - \hat{\mathbf{b}}$ be the vector of changes in individuals' standalone marginal benefits at the optimal intervention. Furthermore, let

$$\alpha_\ell = \frac{1}{(1 - \beta\lambda_\ell)^2}$$

and note that $a_\ell^* = \sqrt{\alpha_\ell b_\ell}$ is the equilibrium action in the ℓ^{th} principal component of \mathbf{G} (see equation (4)).

Theorem 1. Suppose Assumptions 1–3 hold and the network game satisfies Property A. At the optimal intervention, the similarity between \mathbf{y}^* and principal component $\mathbf{u}^\ell(\mathbf{G})$ satisfies the following proportionality:

$$\rho(\mathbf{y}^*, \mathbf{u}^\ell(\mathbf{G})) \propto \rho(\hat{\mathbf{b}}, \mathbf{u}^\ell(\mathbf{G})) \frac{w\alpha_\ell}{\mu - w\alpha_\ell}, \quad \ell = 1, 2, \dots, n, \quad (5)$$

where μ , the shadow price of the planner's budget, is uniquely determined as the solution to

$$\sum_{\ell=1}^n \left(\frac{w\alpha_\ell}{\mu - w\alpha_\ell} \right)^2 \hat{b}_\ell^2 = C \quad (6)$$

and satisfies $\mu > w\alpha_\ell$ for all ℓ , so that all denominators are positive.

We briefly sketch the main argument here and interpret the quantities in the formula. Define $x_\ell = (b_\ell - \hat{b}_\ell)/\hat{b}_\ell$ as the change of b_ℓ , relative to \hat{b}_ℓ . By rewriting the principal's objective and budget constraints in terms of principal components and plugging in the equilibrium

¹⁴In the local public good application (recall Example 2) $w = -1$, and so when $C \geq \|\hat{\mathbf{b}}\|$, the optimal intervention satisfies $b_i^* = 0$. Recalling our change of variables there ($b_i = [\tau - \tilde{b}]/2$), the optimal intervention in that case is to modify the endowment of each individual so that everyone accesses the optimal level of the local public good without investing personally.

condition (4), we can rewrite the maximization problem as

$$\max_{\mathbf{x}} W(\mathbf{b}, \mathbf{G}) = \sum_{\ell=1}^n w\alpha_{\ell}(1+x_{\ell})^2 \hat{b}_{\ell}^2 \quad \text{s.t.} \quad \sum_{\ell=1}^n \hat{b}_{\ell}^2 x_{\ell}^2 \leq C.$$

If the planner allocates a marginal unit of the budget to principal component ℓ , the condition for equality of the marginal return and marginal cost (recalling that μ is the multiplier on the budget constraint) is

$$\underbrace{2\hat{b}_{\ell}^2 \cdot w\alpha_{\ell}(1+x_{\ell})}_{\text{marginal return}} = \underbrace{2\hat{b}_{\ell}^2 \cdot \mu x_{\ell}}_{\text{marginal cost}}.$$

It follows that $\frac{w\alpha_{\ell}}{\mu-w\alpha_{\ell}}$ is exactly the value of x_{ℓ} at which the marginal return and the marginal cost are equalized.¹⁵ Rewriting x_{ℓ} in terms of cosine similarity, that equality implies

$$\frac{w\alpha_{\ell}}{\mu-w\alpha_{\ell}} = x_{\ell}^* = \frac{\|\mathbf{y}^*\| \rho(\mathbf{y}^*, \mathbf{u}^{\ell}(\mathbf{G}))}{\|\hat{\mathbf{b}}\| \rho(\hat{\mathbf{b}}, \mathbf{u}^{\ell}(\mathbf{G}))}.$$

Rearranging this yields the proportionality expression (5) in the theorem. The Lagrange multiplier μ is determined by solving the simple equation (6). Now, given μ , the similarities $\rho(\mathbf{y}^*, \mathbf{u}^{\ell}(\mathbf{G}))$ determine the direction of the optimal intervention \mathbf{y}^* . The magnitude of the intervention is found by exhausting the budget. Thus Theorem 1 provides a full characterization of the optimal intervention.

Next, we discuss the formula for the similarities given in expression (5). The similarity between \mathbf{y}^* and $\mathbf{u}^{\ell}(\mathbf{G})$ measures the extent to which principal component $\mathbf{u}^{\ell}(\mathbf{G})$ is represented in the optimal intervention \mathbf{y}^* . Equation (5) tells us that this is proportional to two factors. The first factor, $\rho(\hat{\mathbf{b}}, \mathbf{u}^{\ell}(\mathbf{G}))$, is a status quo effect corresponding to the similarity between the ℓ^{th} principal component and the status-quo vector $\hat{\mathbf{b}}$. This factor summarizes how much the initial condition influences the optimal intervention for a given budget. The intuition here is that if a given principal component is strongly represented in the status quo vector of standalone incentives, then – because of the convexity of welfare in the principal component basis – changes in that dimension have a particularly significant effect.

The second factor, $\frac{w\alpha_{\ell}}{\mu-w\alpha_{\ell}}$, is determined by two quantities: the eigenvalue corresponding to $\mathbf{u}^{\ell}(\mathbf{G})$ (via $\alpha_{\ell} = \frac{1}{1-\beta\lambda_{\ell}}$), and the budget C (via the shadow price μ). To focus on this second factor, $\frac{w\alpha_{\ell}}{\mu-w\alpha_{\ell}}$, we define the *similarity ratio* of $\mathbf{u}^{\ell}(\mathbf{G})$ to be the fraction

$$r_{\ell}^* = \frac{\rho(\mathbf{y}^*, \mathbf{u}^{\ell}(\mathbf{G}))}{\rho(\hat{\mathbf{b}}, \mathbf{u}^{\ell}(\mathbf{G}))}.$$

Theorem 1 shows that, as we vary ℓ , the similarity ratio r_{ℓ}^* is proportional to $\frac{w\alpha_{\ell}}{\mu-w\alpha_{\ell}}$. It follows that the similarity ratio is greater, in absolute value, for the principal components ℓ with

¹⁵It can be verified that the ratio for every $\ell \in \{1, \dots, n-1\}$, $x_{\ell}/x_{\ell+1}$ is increasing (decreasing) in β for the case of strategic complements (substitutes): thus the intensity of the strategic interaction shapes the relative importance of different principal components.

greater α_ℓ . Intuitively, those are the components where the intervention makes the largest change relative to the status quo profile of incentives. The ordering of these coefficients corresponds to the eigenvalues in a way that depends on the nature of strategic spillovers:

Corollary 1. Suppose Assumptions 1–3 hold and the network game satisfies Property A. If the game is one of strategic complements ($\beta > 0$), then $|r_\ell^*|$ is decreasing in ℓ ; if the game is one of strategic substitutes ($\beta < 0$), then $|r_\ell^*|$ is increasing in ℓ .

In some problems there may be a nonnegativity constraint on actions, in addition to the constraints in problem (IT). Note that as long as the status quo actions $\hat{\mathbf{b}}$ are positive, this constraint will be respected for all C less than some \hat{C} , and so our approach will give information about the relative effects on various components for interventions that are not too large.

4.1. Small and large budgets. The optimal intervention takes especially simple forms in the cases of small and large budgets. From equation (6), we can deduce that the shadow price μ is decreasing in C . For $w > 0$, it follows that an increase in C raises $\frac{w\alpha_\ell}{\mu - w\alpha_\ell}$ and that the principal components with larger α_ℓ become larger in relative terms as well; in other words, if $w > 0$ and $\alpha_\ell > \alpha_{\ell'}$, then $r_\ell^*/r_{\ell'}^*$ is increasing in C .¹⁶ For simplicity of exposition, we suppress the dependence of outcomes on C in the following statement, but note that \mathbf{y}^* and thus the r_ℓ^* are all functions of C .

Proposition 1. Suppose Assumptions 1–3 hold and the network game satisfies Property A. Then the following hold:

1. As $C \rightarrow 0$, in the optimal intervention, $\frac{r_\ell^*}{r_{\ell'}^*} \rightarrow \frac{\alpha_\ell}{\alpha_{\ell'}}$.
2. As $C \rightarrow \infty$, in the optimal intervention
 - 2a. If the game has the strategic complements property, $\beta > 0$, then the similarity of \mathbf{y}^* and the first principal component of the network tends to 1, $\rho(\mathbf{y}^*, \mathbf{u}^1(\mathbf{G})) \rightarrow 1$.
 - 2b. If the game has the strategic substitutes property, $\beta < 0$, then the similarity of \mathbf{y}^* and the last principal component of the network tends to 1, $\rho(\mathbf{y}^*, \mathbf{u}^n(\mathbf{G})) \rightarrow 1$.

This result can be understood by recalling equation (5) in Theorem 1. First, consider the case of small C . When the planner’s budget becomes small, the shadow price μ tends to ∞ .¹⁷ Equation (5) then implies that the similarity ratio of the ℓ^{th} principal component becomes proportional to α_ℓ . Turning now to the case where C grows large, the shadow price converges to $w\alpha_1$ if $\beta > 0$, and to $w\alpha_n$ if $\beta < 0$ (by equation (6)). Plugging this into equation (5), we find that in the case of strategic complements, the optimal intervention shifts individuals’ standalone marginal returns (very nearly) in proportion to the first principal component

¹⁶Analogously, when $w < 0$, $\frac{w\alpha_\ell}{\mu - w\alpha_\ell}$ and $r_\ell^*/r_{\ell'}^*$ are both decreasing in C .

¹⁷As costs are quadratic, small relaxation in the budget around zero can have a large impact on aggregate welfare.

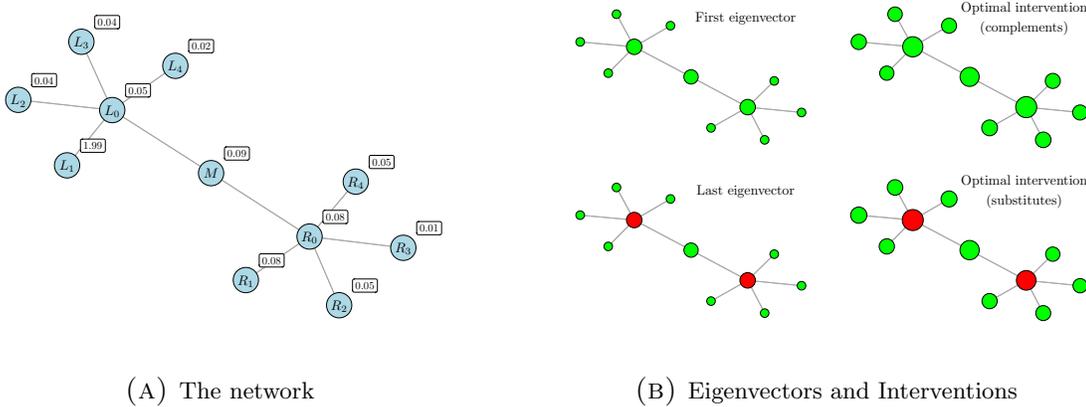


FIGURE 2. Optimal targets with large budgets

of \mathbf{G} , so that $\mathbf{y}^* \rightarrow \sqrt{C}\mathbf{u}^1(\mathbf{G})$. In the case of strategic substitutes, on the other hand, the planner changes individuals' standalone marginal returns (very nearly) in proportion to the last principal component, namely $\mathbf{y}^* \rightarrow \sqrt{C}\mathbf{u}^n(\mathbf{G})$.¹⁸

Figure 2 presents optimal targets when the budget is large – in particular, for $C = 500$. We consider an 11-node undirected network with binary links containing two hubs, L_0 and R_0 , that are connected by an intermediate node M ; the network is depicted in Figure 2(A). The numbers next to the nodes are the status quo standalone marginal returns. Payoffs are as in Example 1. For the case of strategic complements, we set $\beta = 0.1$, and for strategic substitutes we set $\beta = -0.1$. Assumptions 1 and 2 are satisfied and Property A holds. The top-left of Figure 2(B) illustrates the first eigenvector, and the top-right depicts optimal targets in a game with strategic complements. The bottom-left of Figure 2(B) illustrates the last eigenvector, and the bottom-right depicts the optimal targets when the game has strategic substitutes. The node size represents the size of the intervention, $|b_i^* - \hat{b}_i|$; its color represents the sign of the intervention (with green signifying a positive intervention and red indicating a negative intervention).

In line with part 2 of Proposition 1, for large C , the optimal intervention is guided by the “main” component of the network (corresponding to the largest or smallest eigenvalue). Under strategic complements, this is the first eigenvector of the network, which corresponds to individuals' eigenvector centrality.¹⁹ Intuitively, by increasing the standalone marginal return of each individual in proportion to his eigenvector centrality, the planner targets the

¹⁸When individuals' initial standalone marginal returns are zero ($\hat{\mathbf{b}} = \mathbf{0}$), we can dispense with the approximations invoked for a large budget C . Assuming that \mathbf{G} is generic, if $\hat{\mathbf{b}} = \mathbf{0}$ then, regardless of the level of C , the entire budget is spent either (i) on changing b_1 (if $\beta > 0$) or (ii) on changing b_n (if $\beta < 0$). To see this, consider the proof of Theorem 1 and set $\hat{\mathbf{b}} = \mathbf{0}$ in the maximization problem (IT-PC), the principal component version of (IT). Note that if the allocation is not extreme, then the effort can be reallocated profitably among the principal components without changing the cost.

¹⁹Online Appendix Section OA2.1 presents a discussion of eigenvector centrality.

individuals in proportion to their global contributions to strategic feedbacks, and this is welfare maximizing.

Under strategic substitutes, optimal targeting is determined by the last eigenvector of the network, corresponding to its smallest eigenvalue. This network component contains information about the local structure of the network: it determines the way to partition the set of nodes into two sets so that most of the links are across individuals in different sets.²⁰ The optimal intervention increases the standalone marginal returns of all individuals in one set and decreases those of individuals in the other set. The planner wishes to target neighboring nodes asymmetrically, as this reduces possible crowding-out effects that occur due to the strategic substitutes property.

4.2. When are interventions simple? We have just seen examples illustrating how, with large budgets, the intervention is simple: approximately proportional to just one principal component – the top or bottom one. After defining simplicity formally, our final result in this section characterizes how large the budget must be for this approximation to be a close one.

Definition 2 (Simple interventions). An intervention is *simple* if, for all $i \in \mathcal{N}$,

- $b_i - \hat{b}_i = \sqrt{C}u_i^1$ when the game has the strategic complements property ($\beta > 0$),
- $b_i - \hat{b}_i = \sqrt{C}u_i^n$ when the game has the strategic substitutes property ($\beta < 0$).

Let W^* be the aggregate utility under the optimal intervention, and let W^s be the aggregate utility under the simple intervention.

Proposition 2. Suppose $w > 0$, Assumptions 1 and 2 hold, and the network game satisfies Property A.

1. If the game has the strategic complements property, $\beta > 0$, then for any $\epsilon > 0$, if $C > \frac{2\|\hat{\mathbf{b}}\|^2}{\epsilon} \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2$, then $W^*/W^s < 1 + \epsilon$ and $\rho(\mathbf{y}^*, \sqrt{C}\mathbf{u}^1) > \sqrt{1 - \epsilon}$.
2. If the game has the strategic substitutes property, $\beta < 0$, then for any $\epsilon > 0$, if $C > \frac{2\|\hat{\mathbf{b}}\|^2}{\epsilon} \left(\frac{\alpha_{n-1}}{\alpha_n - \alpha_{n-1}} \right)^2$, then $W^*/W^s < 1 + \epsilon$ and $\rho(\mathbf{y}^*, \sqrt{C}\mathbf{u}^n) > \sqrt{1 - \epsilon}$.

Proposition 2 gives a condition on the size of the budget beyond which (a) simple interventions achieve most of the optimal welfare and (b) the optimal intervention is very similar to the simple intervention. This bound depends on the status quo standalone marginal returns and the structure of the network.

We first discuss the dependence on the status quo benefits. Observe that the first term on the right-hand side of the inequality for C is proportional to $\|\hat{\mathbf{b}}\|$. This inequality is therefore

²⁰The last eigenvector of a graph is useful in determining the bipartiteness of a graph and its chromatic number. 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) relate the last eigenvector to a coloring of the underlying graph, that is, a labeling of nodes by a minimal set of integers such that no neighboring nodes share the same label.

easier to satisfy when the status quo standalone marginal returns are smaller, in the sense of having a smaller norm. The inequality is harder to satisfy when these marginal returns are large and/or heterogeneous.²¹

Next, consider the role of the network. Recall that $\alpha_\ell = (1 - \beta\lambda_\ell)^{-2}$; thus if $\beta > 0$, the term $\alpha_2/(\alpha_1 - \alpha_2)$ of the inequality is large when $\lambda_1 - \lambda_2$, the “spectral gap” of the graph, is small. If $\beta < 0$, then the term $\alpha_{n-1}/(\alpha_{n-1} - \alpha_n)$ is small when the “bottom gap” of the graph, the difference $\lambda_{n-1} - \lambda_n$, is small.

We now examine what network features affect these gaps, and illustrate with examples, depicted in Figure 3. The obstacle to simplicity is a strong dependence on the status quo standalone marginal benefits. This dependence will be strong when two different principal components in the network offer similar amplification (all else equal) of interventions in that component. Which of these principal components receives the planner’s focus will depend strongly on the status quo. In such networks, interventions will *not* be simple for reasonable budgets. The merit of Proposition 2 is to show that small spectral or small bottom gap capture this property of the network. Figure 3 illustrates the role of the network structure in shaping the rate (in terms of the size of the budget C) at which the optimal intervention converges to a simple intervention as we vary C . Under strategic complements, the optimal intervention converges to a simple one faster (as we vary C) in a network that has a large spectral gap. Under strategic substitutes, the optimal intervention converges to a simple one faster (as we vary C) in a network that has a large bottom gap.

We now describe which network properties, at a more intuitive level, correspond to having small and large spectral gaps. First, consider the case of strategic complements. A standard fact is that the two largest eigenvalues can be expressed in terms of the corresponding eigenvectors as follows:

$$\lambda_1 = \max_{\mathbf{u}: \|\mathbf{u}\|=1} \sum_{i,j \in \mathcal{N}} g_{ij} u_i u_j \quad \lambda_2 = \max_{\substack{\mathbf{u}: \|\mathbf{u}\|=1 \\ \mathbf{u} \cdot \mathbf{u}^1 = 0}} \sum_{i,j \in \mathcal{N}} g_{ij} u_i u_j.$$

Eigenvector $\mathbf{u}^1 = \arg \max_{\mathbf{u}: \|\mathbf{u}\|=1} \sum_{i,j \in \mathcal{N}} g_{ij} u_i u_j$ (corresponding to λ_1) assigns the same sign – say, positive to all nodes in the network. Clearly, eigenvector \mathbf{u}^2 must assign negative values to some of the nodes (as it is orthogonal to \mathbf{u}^1). In the network on the left side of Figure 4(A) any such assignment will result in many adjacent nodes having opposite sign entries of \mathbf{u}^2 ; as a result, many terms in the expression for λ_2 will be negative, and λ_2 will be much smaller than λ_1 , leading to a large spectral gap. In the network on the right side of 4(A), \mathbf{u}^2 will have positive-sign entries to nodes in one community and negative-sign entries for nodes in the other community. Because there are few edges between the communities, λ_2 turns out to be almost as large as λ_1 . This will yield a small spectral gap. Thus, spectral

²¹Recall that $\|\frac{1}{n}\hat{\mathbf{b}}\|^2$ is equal to the sum of $\left(\frac{1}{n}\sum_{i \in \mathcal{N}} \hat{b}_i\right)^2$ (the squared mean of the entries of \mathbf{b}) and the sum of squared deviations of the entries of the vector $\hat{\mathbf{b}}$ from their mean.

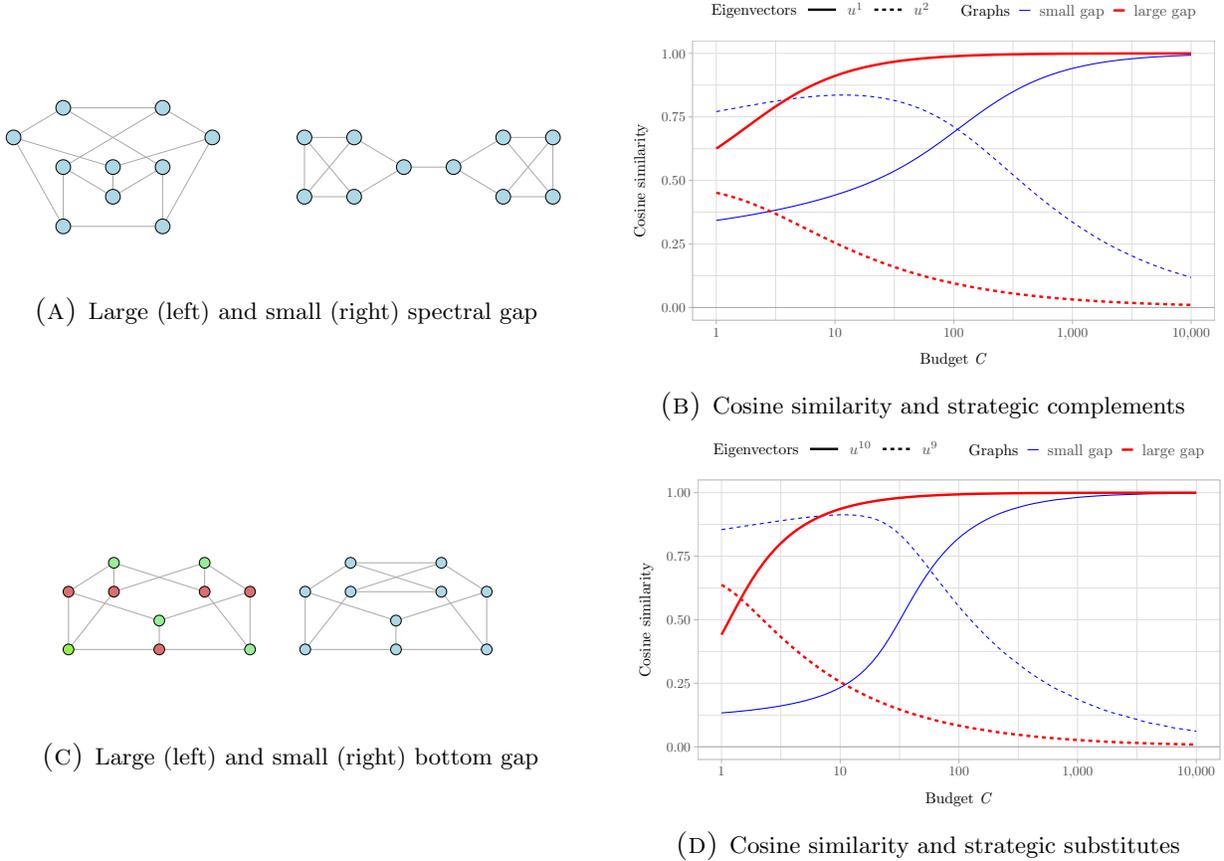


FIGURE 3. Spectral gap, bottom gap, and optimal interventions

gap measures the level of “cohesiveness” of the network, and it is this property that dictates fast convergence to simple interventions.²²

Turning next to strategic substitutes, recall that the smallest two eigenvalues, λ_n and λ_{n-1} , can be written in terms of the corresponding eigenvectors as follows:

$$\lambda_n = \min_{\mathbf{u}: \|\mathbf{u}\|=1} \sum_{i,j \in \mathcal{N}} g_{ij} u_i u_j \quad \lambda_{n-1} = \min_{\substack{\mathbf{u}: \|\mathbf{u}\|=1 \\ \mathbf{u} \cdot \mathbf{u}^n = 0}} \sum_{i,j \in \mathcal{N}} g_{ij} u_i u_j. \quad (7)$$

This tells us that λ_n is far to the left on the real number line²³ when the eigenvector $\mathbf{u}^n = \arg \min_{\mathbf{u}: \|\mathbf{u}\|=1} \sum_{i,j \in \mathcal{N}} g_{ij} u_i u_j$ (corresponding to λ_n) assigns opposite signs to most pairs of adjacent nodes. In other words, the last eigenvalue is small when nodes can be partitioned into two sets and most of the connections are across sets, and thus λ_n is minimized in a bipartite graph. The second-smallest eigenvalue of \mathbf{G} reflects the extent to which the next-best eigenvector (orthogonal to \mathbf{u}^n) is good at solving the same minimization problem.

²²See Hartfiel and Meyer (1998), Levin et al. (2009), and Golub and Jackson (2012) for discussions and further citations to the literature on spectral gaps.

²³The eigenvalue is in fact negative, so as a consequence of the assumption that $g_{ii} = 0$ for all i : In this case the trace of \mathbf{G} is zero, and therefore its eigenvalues sum to 0. By the Perron-Frobenius Theorem, the maximum eigenvalue of the nonnegative matrix \mathbf{G} is positive, so the minimum one must be negative.

Hence, the bottom gap of \mathbf{G} is small when there are two orthogonal ways to partition the network into two sets so that, either way, the “quality” of the bipartition, as measured by $\sum_{i,j \in \mathcal{N}} g_{ij} u_i u_j$, is similar.

We illustrate Proposition 2 with a comparison of the two graphs in Figure 4(C). The left-hand graph is bipartite: the last eigenvalue is $\lambda_n = -3$ and the second last eigenvalue is $\lambda_{n-1} = -1.64$. By contrast, in the graph on the right of Figure 4(C), the bottom eigenvalue $\lambda_n = -2.62$, while the second lowest is $\lambda_{n-1} = -2.30$. This yields a much smaller bottom gap.²⁴ This difference in bottom gap is reflected in optimal targeting policy in Figure 4(D): in the graph with large bottom gap, the optimal intervention puts little weight on the eigenvector \mathbf{u}^{n-1} for a relatively small budget; it takes a much larger budget under a small bottom gap.

We conclude by noting the influence of the status quo standalone marginal returns in shaping optimal interventions for small budgets. For a small budget C , the cosine similarity of the optimal intervention for non-main network components can be higher than the one for the main component. This is true when the status quo $\hat{\mathbf{b}}$ is similar to some of the non-main network components; see Figures 4(B) and Figure 4(D).

5. INCOMPLETE INFORMATION

In the basic model, we assumed that the planner knows the standalone marginal returns of every individual. This section extends the analysis to settings where the planner does *not* know these parameters. For ease of exposition, we focus on network games that satisfy Property A.

Formally, fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The planner’s belief over states is given by \mathbb{P} . This represents the planner’s uncertainty, given all her information. The planner has control over the random vector (r.v.) \mathcal{B} , that is, a function $\mathcal{B} : \Omega \rightarrow \mathbb{R}^n$. The choice of \mathcal{B} determines the cost of intervention. A realization of the random vector is denoted by \mathbf{b} . This realization is common knowledge among individuals when they choose their actions. Thus, the game individuals play is one of complete information.²⁵ We also define a function K that gives the cost $K(\mathcal{B})$ of implementing the random variable \mathcal{B} .²⁶

We solve the following incomplete-information intervention problem:

$$\begin{aligned} \text{choose r.v. } \mathcal{B} \text{ to maximize } & \mathbb{E}[W(\mathbf{b}; \mathbf{G})] & (\text{IT-G}) \\ \text{s.t. } & [\mathbf{I} - \beta \mathbf{G}] \mathbf{a}^* = \mathbf{b}, \\ & K(\mathcal{B}) \leq C. \end{aligned}$$

²⁴Intuitively, because \mathbf{u}^n does not correspond to a perfect bipartition, it is easier for a vector orthogonal to \mathbf{u}^n to achieve a similarly low value of $\sum_{i,j \in \mathcal{N}} g_{ij} u_i u_j$.

²⁵It is possible to go further and allow for incomplete information among the individuals about each other’s b_i . We do not pursue this substantial generalization here; see Golub and Morris (2017) and Lambert et al. (2018) for analyses in this direction.

²⁶The domain of this function is the set of all random vectors taking values in \mathbb{R}^n defined on our probability space.

Note that the intervention problem (IT) under complete information is a special case of a degenerate r.v. \mathcal{B} : one in which the planner knows the vector of standalone marginal returns exactly and implements a deterministic adjustment relative to it.

To guide our modeling of the cost of intervention, we now review the features of the distribution of \mathcal{B} that matter for aggregate welfare. For network games that satisfy Property A, we can write:

$$\mathbb{E}[W(\mathbf{b}; \mathbf{G})] = w\mathbb{E}[(\mathbf{a}^*)^\top \mathbf{a}^*] = w\mathbb{E}[\underline{\mathbf{a}}^\top \underline{\mathbf{a}}] = w \sum_{\ell=1}^n \alpha_\ell (\mathbb{E}[b_\ell]^2 + \text{Var}[b_\ell]). \quad (8)$$

In words, welfare is determined by the mean and variance of the realized components b_ℓ ; these in turn are determined by the first and second moments of the chosen random variable \mathcal{B} . In view of this, we will consider intervention problems where the planner can modify the mean and the covariance matrix of \mathcal{B} , and the cost of intervention depends only on these modifications.

5.1. Mean shifts. We first consider an intervention where there is an arbitrarily distributed vector of standalone marginal returns and the planner's intervention shifts it in a deterministic way. Formally, fix a random variable $\hat{\mathcal{B}}$, called the status quo, with typical realization $\hat{\mathbf{b}}$. The planner's policy is given by $\mathbf{b} = \hat{\mathbf{b}} + \mathbf{y}$, where $\mathbf{y} \in \mathbb{R}^n$ is a deterministic vector. We denote the corresponding random variable by $\mathcal{B}_\mathbf{y}$. In terms of interpretation, note that implementing this policy does not require knowing $\hat{\mathbf{b}}$ as long as the planner has an instrument that shifts incentives.

Assumption 4. The cost of implementing r.v. $\mathcal{B}_\mathbf{y}$ is

$$K(\mathcal{B}_\mathbf{y}) = \sum_{i \in \mathcal{N}} y_i^2,$$

and $K(\mathcal{B})$ is ∞ for any other random variable.

In contrast to the analysis of Theorem 1, the vector $\hat{\mathbf{b}}$ is a random variable. But we take the analogue of the cost function used there, noting that in the deterministic setting this formula held with $\mathbf{y} = \mathbf{b} - \hat{\mathbf{b}}$.

Proposition 3. Consider problem (IT-G) with the cost of intervention satisfying Assumption 4. Suppose Assumptions 1 and 2 hold and the network game satisfies Property A. The optimal intervention policy \mathcal{B}^* is equal to $\mathcal{B}_{\mathbf{y}^*}$, where \mathbf{y}^* is the optimal intervention in the deterministic problem with $\bar{\mathbf{b}} = \mathbb{E}[\hat{\mathbf{b}}]$ as the status quo vector of standalone marginal returns.

5.2. Intervention on variances. We next consider the case where the planner faces a vector of means, fixed at $\bar{\mathbf{b}}$, and, subject to that, can choose any random variable \mathcal{B} . The difference in the expected welfare for two different interventions \mathcal{B} and $\tilde{\mathcal{B}}$ depends only on the variance-covariance matrix of \mathcal{B} and $\tilde{\mathcal{B}}$. Thus, the planner effectively faces *the problem*

of intervening on variances. We prove a result on optimal intervention for all cost functions satisfying certain symmetries.

Assumption 5. The cost function satisfies two properties: (a) $K(\mathcal{B}) = \infty$ if $\mathbb{E}\mathbf{b} \neq \bar{\mathbf{b}}$; (b) $K(\mathcal{B}) = K(\tilde{\mathcal{B}})$ if $\tilde{\mathbf{b}} - \bar{\mathbf{b}} = \mathbf{O}(\mathbf{b} - \bar{\mathbf{b}})$, where \mathbf{O} is an orthogonal matrix. Analogous to our other notation, we use $\tilde{\mathbf{b}}$ for realizations of the random vector with distribution $\tilde{\mathcal{B}}$.

Part (a) is a restriction on feasible interventions, namely a restriction to interventions that are mean-neutral. Part (b) means that rotations of coordinates around the mean do not affect the cost of implementing a given distribution. This assumption gives the cost a directional neutrality, which ensures that our results are driven by the benefits side rather than by asymmetries operating through the costs. For example, let $\Sigma_{\mathcal{B}}$ be the variance-covariance matrix of the random variable \mathcal{B} . In particular, $\sigma_{ii}^{\mathcal{B}}$ is the variance of b_i . Suppose that the cost of implementing \mathcal{B} with $\mathbb{E}\mathbf{b} = \bar{\mathbf{b}}$ is a function of the sum of the variances of the b_i :

$$K(\mathcal{B}) = \begin{cases} \phi\left(\sum_{i \in \mathcal{N}} \sigma_{ii}^{\mathcal{B}}\right) & \text{if } \mathbb{E}\mathbf{b} = \bar{\mathbf{b}} \\ \infty & \text{otherwise.} \end{cases} \quad (9)$$

The cost function (9) satisfies property (a) of Assumption 5. Moreover, it satisfies property (b) of Assumption 5 because $\sum_{i \in \mathcal{N}} \sigma_{ii}^{\mathcal{B}} = \text{trace } \Sigma_{\mathcal{B}}$; this trace is the sum of the eigenvalues of $\Sigma_{\mathcal{B}}$, which is invariant to the transformation defined in (b).²⁷

Proposition 4 (Variance control). Consider problem (IT-G) with the cost of intervention satisfying Assumption 5. Suppose Assumptions 1 and 2 hold and the network game satisfies Property A. Let the optimal intervention be \mathcal{B}^* . We have the following:

1. Suppose the planner likes variance (i.e., $w > 0$). If the game has strategic complements ($\beta > 0$), then $\text{Var}(\mathbf{u}^{\ell}(\mathbf{G}) \cdot \mathbf{b}^*)$ is weakly decreasing in ℓ ; if the game has strategic substitutes ($\beta < 0$), then $\text{Var}(\mathbf{u}^{\ell}(\mathbf{G}) \cdot \mathbf{b}^*)$ is weakly increasing in ℓ .
2. Suppose the planner dislikes variance (i.e., $w < 0$). If the game has strategic complements ($\beta > 0$), then $\text{Var}(\mathbf{u}^{\ell}(\mathbf{G}) \cdot \mathbf{b}^*)$ is weakly increasing in ℓ ; if the game has strategic substitutes ($\beta < 0$), then $\text{Var}(\mathbf{u}^{\ell}(\mathbf{G}) \cdot \mathbf{b}^*)$ is weakly decreasing in ℓ .

We now provide the intuition for Proposition 4. Shocks to individuals' standalone marginal returns create variability in the players' equilibrium actions. The assumption that the intervention is mean neutral (part (a) of Assumption 5) leaves the planner to control only the variances and covariances of these marginal returns with her intervention. Hence, the solution to the intervention problem describes what the planner should do to induce volatilities in actions that maximize the ex-ante expected welfare.

Suppose first that investments are strategic complements. Then a perfectly correlated shock in individual standalone marginal returns is amplified by strategic interaction. In

²⁷When we look at the variance-covariance matrix of $\tilde{\mathbf{b}}$ defined by $\tilde{\mathbf{b}} - \bar{\mathbf{b}} = \mathbf{O}(\mathbf{b} - \bar{\mathbf{b}})$, the variance-covariance matrix becomes $\mathbf{O}\Sigma\mathbf{O}^{\top}$, and this has the same eigenvalues and therefore the same trace.

fact, the type of shock that is most amplifying (at a given size) is the one that is perfectly correlated across individuals, with the magnitude of a given individual's shock proportional to the first principal component (his eigenvector centrality). These shocks are exactly what $b_1^* = \mathbf{u}^1(\mathbf{G}) \cdot \mathbf{b}^*$ captures. Hence, this is the dimension of volatility that the planner most wants to increase if she likes variance in actions ($w > 0$) and most wants to decrease if she dislikes variance in actions ($w < 0$).

If investments are strategic substitutes, then a perfectly correlated shock does not create a lot of variance in actions: The first-order response of all individuals to an increase in their standalone marginal returns is to increase investment, but that in turn makes all individuals decrease their investment somewhat because of the strategic substitutability with their neighbors. Hence, highly positively correlated shocks do not translate into high volatility. The shock profiles that create most variability in actions are the ones in which neighbors have *negatively* correlated shocks. A planner that likes variability in actions will then prioritize such shocks. Because the last eigenvector of the system is correlated with those shocks that have opposite effects on neighbors, this is exactly the type of volatility that is of greatest concern, and this is what the planner will focus on most.

Example 3 (Illustration in the case of the circle). Figure 1 depicts six of the eigenvectors/principal components of a circle network with 14 nodes. The first principal component is a positive vector and so \mathcal{B} projected on $\mathbf{u}^1(\mathbf{G})$ captures positively correlated shocks across all players. The second principal component (top left panel of Figure 1) splits the graph into two sides, one with positive entries and the other with negative entries. Hence, \mathcal{B} projected on $\mathbf{u}^2(\mathbf{G})$ captures shocks that are highly positively correlated on each side of the circle network, with the two opposite sides of the circle being anti-correlated. As we move along the sequence, we can see that \mathcal{B} projected on the ℓ^{th} eigenvector represents shocks that are more and more local. At the extreme, \mathcal{B} projected on $\mathbf{u}^{14}(\mathbf{G})$, (bottom-right panel of Figure 1) captures the component of shocks that is perfectly anti-correlated across neighbors.²⁸

6. CONCLUDING REMARKS

We study the problem of a planner who seeks to optimally target incentive changes in a network game. Our framework allows for a broad class of strategic and non-strategic spillovers across neighbors. The main contribution of the paper is methodological: we show that principal components of the network of interaction provide a useful basis for analyzing the effects of an intervention. This decomposition leads to our main result: there is a close relation between the nature of the game (complements or substitutes) and the weight that different principal components receive in the optimal intervention. To develop these ideas in

²⁸As usual in this example, a generic \mathbf{G} will not be perfectly symmetric and so a particular orientation of these eigenvectors will be selected.

the simplest way, we have focused on a model in which the matrix of interaction is symmetric, the costs of intervention are quadratic, and the intervention itself takes the form of altering the standalone benefits. In the Online Appendix we relax these restrictions and develop extensions of our approach to non-symmetric matrices of interaction, to more general costs of intervention, and to environments where interventions occur via monetary incentives for activity. We also relax Property A, a technical condition which facilitated our basic analysis, and cover a more general class of externalities.

We briefly mention two further applications. In some circumstances, the planner seeks a budget-balanced tax/subsidy scheme in order to improve the economic outcome. In an oligopoly market, for example, a planner could tax some suppliers, thereby increasing their marginal costs, and then use that tax revenue to subsidize other suppliers. The planner will solve a problem similar to the one we have studied here, with the important difference that she will face a different constraint, namely, a budget-balance constraint. In ongoing work, Galeotti et al. (2018) show that the principal component approach 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.

We have focused on interventions that alter the standalone marginal returns of individuals. Another interesting problem is the study of interventions that alter the matrix of interaction. We hope this paper stimulates further work along these lines.

REFERENCES

- AKBARPOUR, M., S. MALLADI, AND A. SABERI (2017): “Diffusion, Seeding, and the Value of Network Information,” Available at SSRN: <https://ssrn.com/abstract=3062830>.
- ALLOUCH, N. (2015): “On the private provision of public goods on networks,” *Journal of Economic Theory*, 157, 527–552.
- (2017): “Aggregation in Networks,” University of Kent, School of Economics Discussion Paper 1718.
- 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 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. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): “The Diffusion of Microfinance,” *Science*, 341, 1236498.
- BELHAJ, M. AND F. DEROIAN (2017): “Contracting on Networks,” Mimeo.
- 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 (2007): “Public Goods in Networks,” *Journal of Economic Theory*, 135, 478–494.
- 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.
- CHUNG, F. R. AND F. C. GRAHAM (1997): *Spectral graph theory*, 92, American Mathematical Soc.
- CVETKOVIC, D., D. M. CVETKOVIĆ, P. ROWLINSON, AND S. SIMIC (1997): *Eigenspaces of graphs*, 66, Cambridge University Press.
- DAVIES, E. B., G. M. GLADWELL, J. LEYDOLD, AND P. F. STADLER (2001): “Discrete nodal domain theorems,” *Linear Algebra and its Applications*, 336, 51 – 60.
- DEMANGE, G. (2017): “Optimal Targeting Strategies in a Network under Complementarities,” *Games and Economic Behaviour*, 105, 84–103.
- DESAI, M. AND V. RAO (1994): “A Characterization of the Smallest Eigenvalue of a Graph,” *Journal of Graph Theory*, 18, 181–194.

- FAINMESSER, I. AND A. GALEOTTI (2017): “Pricing Network Effects,” *Review of Economic Studies*, 83, 165–198.
- 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.
- (2010): “The Law of the Few,” *American Economic Review*, 100, 1468–92.
- GALEOTTI, A., S. GOYAL, M. JACKSON, F. VEGA-REDONDO, AND L. YARIV (2010): “Network Games,” *Review of Economic Studies*, 77, 218–244.
- 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 (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): “Expectations, Networks and Conventions,” Available at SSRN: <https://ssrn.com/abstract=2979086>.
- GOYAL, S., J. MORAGA, AND M. VAN DER LEIJ (2006): “Economics: an emerging small world?” *Journal of Political Economy*, 114, 403–412.
- 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.
- JACKSON, M., B. W. ROGERS, AND Y. ZENOU (2017): “The Economic Consequences of Social-Network Structure,” *Journal of Economic Literature*, 55, 49–95.
- 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*.
- LAMBERT, N. S., G. MARTINI, AND M. OSTROVSKY (2018): “Quadratic Games,” Available at NBER: <https://www.nber.org/papers/w24914>.
- LEDUC, M. V., M. O. JACKSON, AND R. JOHARI (2017): “Pricing and referrals in diffusion on networks,” *Games and Economic Behavior*, 104, 568–594.
- LEVIN, D. A., Y. PERES, AND E. L. WILMER (2009): *Markov Chains and Mixing Times*, Providence, RI: American Mathematical Society.
- MORRIS, S. AND H. S. SHIN (2002): “Social Value of Public Information,” *American Economic Review*, 92, 1521–1534.
- ROGERS, E. (1983): *Diffusion of Innovations. Third Edition*, Free Press. New York.
- SPIELMAN, D. A. (2007): “Spectral graph theory and its applications,” in *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07)*, IEEE, 29–38.

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.

APPENDIX A. PROOFS

Proof of Theorem 1. We wish to solve

$$\begin{aligned} \max_{\mathbf{b}} \quad & w \mathbf{a}^\top \mathbf{a} \\ \text{s.t.} \quad & [\mathbf{I} - \beta \mathbf{G}] \mathbf{a}^* = \mathbf{b}, \\ & \sum_{i \in \mathcal{N}} (b_i - \hat{b}_i)^2 \leq C. \end{aligned}$$

We transform the maximization problem into the basis given by the principal components of \mathbf{G} . To this end, we first rewrite the cost and the objective in the principal components basis, using the fact that norms do not change under the orthogonal transformation \mathbf{U}^\top . (The norm symbol $\|\cdot\|$ always refers to the Euclidean norm.) Letting $\mathbf{y} = \mathbf{b} - \hat{\mathbf{b}}$,

$$K(\mathbf{b}, \hat{\mathbf{b}}) = \sum_{i \in \mathcal{N}} y_i^2 = \|\mathbf{y}\|_2^2 = \sum_{\ell=1}^n \underline{y}_\ell^2$$

and

$$w \mathbf{a}^\top \mathbf{a} = w \|\mathbf{a}\|^2 = w \|\underline{\mathbf{a}}\|^2 = w \underline{\mathbf{a}}^\top \underline{\mathbf{a}}.$$

By recalling that, in equilibrium, $\underline{\mathbf{a}}^* = [\mathbf{I} - \beta \mathbf{\Lambda}]^{-1} \underline{\mathbf{b}}$, and using the definition $\alpha_\ell = \frac{1}{(1 - \beta \lambda_\ell(\mathbf{G}))^2}$, the intervention problem (IT) can be rewritten as:

$$\begin{aligned} \max_{\underline{\mathbf{b}}} \quad & w \sum_{\ell=1}^n \alpha_\ell \underline{b}_\ell^2 && \text{(IT-PC)} \\ \text{s.t.} \quad & \sum_{\ell=1}^n \underline{y}_\ell^2 \leq C. \end{aligned}$$

We now transform the problem so that the control variable is \mathbf{x} where $x_\ell = y_\ell / \hat{b}_\ell$. We obtain

$$\begin{aligned} \max_{\mathbf{x}} \quad & w \sum_{\ell=1}^n \ell = 1^n \alpha_\ell (1 + x_\ell)^2 \hat{b}_\ell^2 \\ \text{s.t.} \quad & \sum_{\ell=1}^n \hat{b}_\ell^2 x_\ell^2 \leq C \end{aligned}$$

Note that, for all ℓ , α_ℓ are well-defined (by Assumption 1) and strictly positive (by genericity of \mathbf{G}). This has two implications.²⁹

²⁹Note that if Assumption 3 does not hold (that is, $w < 0$ and $\sum_{\ell=1}^n \hat{b}_\ell^2 \leq C$) then the optimal solution is $x_\ell^* = -1$ for all ℓ . This is what we ruled out with Assumption 3, before Theorem 1.

First, at the optimal solution \mathbf{x}^* the resource constraint problem must bind. To see this, note that Assumption 3 says that either $w > 0$, or $w < 0$ and $\sum_{\ell=1}^n \hat{b}_\ell^2 > C$. Suppose that at the optimal solution the constraint does not bind. Then, without violating the constraint, we can slightly increase or decrease any x_ℓ . If $w > 0$ (resp. $w < 0$) the increase or the decrease is guaranteed to increase (resp. decrease) the corresponding $(x_\ell + 1)^2$ (since the α_ℓ are all strictly positive).

Second, we show that the optimal solution \mathbf{x}^* satisfies $x_\ell^* \geq 0$ for every ℓ if $w > 0$, and $x_\ell^* \in [-1, 0]$ for every ℓ if $w < 0$. Suppose $w > 0$ and, for some ℓ , $x_\ell^* < 0$. Then $[-x_\ell^* + 1]^2 > [x_\ell^* + 1]^2$. Since $w > 0$ and every α_ℓ is positive, we can raise the aggregate utility without changing the cost by flipping the sign of x_ℓ^* . Analogously, suppose $w < 0$. It is clear that if $x_\ell^* < -1$, then by setting $x_\ell = -1$ the objective improves and the constraint is relaxed; hence, at the optimum, $x_\ell^* \geq -1$. Suppose next that $x_\ell > 0$ for some ℓ . Then $[-x_\ell^* + 1]^2 < [x_\ell^* + 1]^2$. Since $w < 0$ and every α_ℓ is positive, we can improve the value of the objective function without changing the cost by flipping the sign of x_ℓ^* .

We now complete the proof. Observe that the Lagrangian corresponding to the maximization problem is

$$\mathcal{L} = w \sum_{\ell=1}^n \alpha_\ell (1 + x_\ell)^2 \hat{b}_\ell + \mu \left[C - \sum_{\ell=1}^n \hat{b}_\ell^2 x_\ell^2 \right].$$

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

$$0 = \frac{\partial \mathcal{L}}{\partial x_\ell} = 2\hat{b}_\ell^2 [w\alpha_\ell(1 + x_\ell^*) - \mu x_\ell^*] = 0. \quad (10)$$

We take a generic $\hat{\mathbf{b}}$ such that $\hat{b}_\ell \neq 0$ for all ℓ . If for some ℓ we had $\mu = w\alpha_\ell$ then the right-hand side of the second equality in (10) would be $2\hat{b}_\ell^2 w\alpha_\ell$, which, by the generic assumption we just made and the positivity of α_ℓ , would contradict (10). Thus, the following holds with a nonzero denominator:

$$x_\ell^* = \frac{w\alpha_\ell}{\mu - w\alpha_\ell},$$

and the Lagrange multiplier μ is therefore pinned down by

$$\sum_{\ell=1}^n w^2 \hat{b}_\ell^2 \left(\frac{\alpha_\ell}{\mu - w\alpha_\ell} \right)^2 = C.$$

Note finally that

$$\rho(\mathbf{y}^*, \mathbf{u}^\ell(\mathbf{G})) = \frac{\mathbf{y}^* \cdot \mathbf{u}^\ell(\mathbf{G})}{\|\mathbf{y}^*\| \|\mathbf{u}^\ell(\mathbf{G})\|} = \frac{y_\ell^*}{\sqrt{C}} = \frac{\hat{b}_\ell x_\ell^*}{\sqrt{C}} = \frac{\|\hat{\mathbf{b}}\|}{\sqrt{C}} \rho(\hat{\mathbf{b}}, \mathbf{u}^\ell(\mathbf{G}))_{x_\ell^*} \propto \rho(\hat{\mathbf{b}}, \mathbf{u}^\ell(\mathbf{G}))_{x_\ell^*}.$$

□

Proof of Proposition 1. Part 1. From expression 6 of Theorem 1, it follows that if $C \rightarrow 0$ then $\mu \rightarrow \infty$. The result follows by noticing that

$$\frac{r_\ell^*}{r_{\ell'}^*} = \frac{\alpha_\ell \mu - w\alpha_\ell'}{\alpha_{\ell'} \mu - w\alpha_{\ell'}}.$$

Part 2. Suppose that $\beta > 0$. Using the derivation of the last part of the proof of Theorem 1, we write:

$$\rho(\mathbf{y}^*, \mathbf{u}^\ell(\mathbf{G})) = \frac{\|\hat{\mathbf{b}}\|}{\sqrt{C}} \rho(\hat{\mathbf{b}}, \mathbf{u}^\ell(\mathbf{G})) x_\ell^*,$$

with $x_\ell^* = \frac{w\alpha_\ell}{\mu - w\alpha_\ell}$. From expression 6 of Theorem 1, it follows that if $C \rightarrow \infty$ then $\mu \rightarrow w\alpha_1$. This implies that $x_\ell^* \rightarrow \frac{\alpha_\ell}{\alpha_1 - \alpha_\ell}$ for all $\ell \neq 1$. As a result, if $C \rightarrow \infty$ then $\rho(\mathbf{y}^*, \mathbf{u}^\ell(\mathbf{G})) \rightarrow 0$ for all $\ell \neq 1$. Furthermore, we can rewrite expression 6 of Theorem 1 as

$$\sum_{\ell=1}^n \left(\|\hat{\mathbf{b}}\| \rho(\hat{\mathbf{b}}, \mathbf{u}^\ell(\mathbf{G})) \frac{x_\ell^*}{\sqrt{C}} \right)^2 = 1,$$

and therefore

$$\lim_{C \rightarrow \infty} \sum_{\ell=1}^n \left(\|\hat{\mathbf{b}}\| \rho(\hat{\mathbf{b}}, \mathbf{u}^\ell(\mathbf{G})) \frac{x_\ell^*}{\sqrt{C}} \right)^2 = \lim_{C \rightarrow \infty} \left(\|\hat{\mathbf{b}}\| \rho(\hat{\mathbf{b}}, \mathbf{u}^1(\mathbf{G})) \frac{x_1^*}{\sqrt{C}} \right)^2 = 1,$$

where the first equality follows because $x_\ell^* \rightarrow \frac{\alpha_\ell}{\alpha_1 - \alpha_\ell}$ for all $\ell \neq 1$. The proof for the case of $\beta < 0$ follows the same steps, with the only exception that if $C \rightarrow \infty$ then $\mu \rightarrow w\alpha_n$. \square

Proof of Proposition 2. We first prove the result on welfare and then turn to the result on cosine similarity.

Welfare. Consider the case of strategic complementarities, $\beta > 0$. Define by $\tilde{\mathbf{x}}$ the simple intervention, and note that $\tilde{x}_1 = \sqrt{C}/\hat{b}_1$ and that $\tilde{x}_\ell = 0$ for all $\ell > 1$. The aggregate utility obtained under the simple intervention is:

$$W^s = \sum_{\ell=1}^n \hat{b}_\ell^2 \alpha_\ell (1 + \tilde{x}_\ell)^2 = \hat{b}_1^2 \alpha_1 \tilde{x}_1 (\tilde{x}_1 + 2) + \sum_{\ell=1}^n \alpha_\ell \hat{b}_\ell^2.$$

The aggregate utility at the optimal intervention is

$$W^* = \sum_{\ell=1}^n \hat{b}_\ell^2 \alpha_\ell (1 + x_\ell^*)^2 = \hat{b}_1^2 \alpha_1 x_1^* (x_1^* + 2) + \sum_{\ell=2}^n \hat{b}_\ell^2 \alpha_\ell x_\ell^* (x_\ell^* + 2) + \sum_{\ell=1}^n \alpha_\ell \hat{b}_\ell^2$$

Hence

$$\begin{aligned} \frac{W^*}{W^s} &= \frac{\hat{b}_1^2 \alpha_1 x_1^* (x_1^* + 2) + \sum_{\ell=1}^n \alpha_\ell \hat{b}_\ell^2}{\hat{b}_1^2 \alpha_1 \tilde{x}_1 (\tilde{x}_1 + 2) + \sum_{\ell=1}^n \alpha_\ell \hat{b}_\ell^2} + \frac{\sum_{\ell=2}^n \hat{b}_\ell^2 \alpha_\ell x_\ell^* (x_\ell^* + 2)}{\hat{b}_1^2 \alpha_1 \tilde{x}_1 (\tilde{x}_1 + 2) + \sum_{\ell=1}^n \alpha_\ell \hat{b}_\ell^2} \\ &\leq 1 + \frac{\sum_{\ell=2}^n \hat{b}_\ell^2 \alpha_\ell x_\ell^* (x_\ell^* + 2)}{\hat{b}_1^2 \alpha_1 \tilde{x}_1 (\tilde{x}_1 + 2) + \sum_{\ell=1}^n \alpha_\ell \hat{b}_\ell^2} \quad \text{as } \tilde{x}_1 \geq x_1^* \end{aligned}$$

$$\begin{aligned}
 &\leq 1 + \frac{\sum_{\ell=2}^n \hat{b}_\ell^2 \alpha_\ell x_\ell^* (x_\ell^* + 2)}{\hat{b}_1^2 \alpha_1 \tilde{x}_1^2} && \text{summands in denominator are positive} \\
 &= 1 + \frac{\sum_{\ell=2}^n \hat{b}_\ell^2 \alpha_\ell x_\ell^* (x_\ell^* + 2)}{\alpha_1 C} && \underline{b}_1^2 \tilde{x}_1^2 = C; \text{ see below} \\
 &\leq 1 + \frac{2\alpha_1 - \alpha_2}{\alpha_1} \frac{\|\hat{\mathbf{b}}\|^2}{C} \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2 && \text{see calculation below} \\
 &\leq 1 + \frac{2\|\hat{\mathbf{b}}\|^2}{C} \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2.
 \end{aligned}$$

The fact $\underline{b}_1^2 \tilde{x}_1^2 = C$, used above, follows because the simple policy allocates the entire budget to changing \underline{b}_1 . The inequality after that statement follows because

$$\begin{aligned}
 \sum_{\ell=2}^n \hat{b}_\ell^2 \alpha_\ell x_\ell^* (x_\ell^* + 2) &\leq \alpha_2 \sum_{\ell=2}^n \hat{b}_\ell^2 x_\ell^* (x_\ell^* + 2) && \text{ordering of the } \alpha_\ell \\
 &\leq \alpha_2 x_2^* (x_2^* + 2) \sum_{\ell=2}^n \hat{b}_\ell^2 && \text{Corollary 1} \\
 &\leq \alpha_2 \frac{w\alpha_2}{\mu - w\alpha_2} \left(\frac{w\alpha_2}{\mu - w\alpha_2} + 2 \right) \sum_{\ell=2}^n \hat{b}_\ell^2 && \text{Theorem 1} \\
 &\leq \alpha_2 \frac{w\alpha_2}{w\alpha_1 - w\alpha_2} \left(\frac{w\alpha_2}{w\alpha_1 - w\alpha_2} + 2 \right) \|\hat{\mathbf{b}}\|^2 \\
 &= \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2 (2\alpha_1 - \alpha_2) \|\hat{\mathbf{b}}\|^2
 \end{aligned}$$

Hence, the inequality

$$C > \frac{2\|\hat{\mathbf{b}}\|^2}{\epsilon} \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2$$

is sufficient to establish that $\frac{W^*}{W_s} < 1 + \epsilon$. The proof for the case of strategic substitutes follows the same steps; the only difference is that we use α_n instead of α_1 and α_{n-1} instead of α_2 .

Cosine similarity. We now turn to the cosine similarity result. We focus on the case of strategic complements. The proof for the case of strategic substitutes is analogous. We start by writing a useful explicit expression for

$$\rho(\Delta \mathbf{b}^*, \sqrt{C} \mathbf{u}^1) = \frac{(\mathbf{b}^* - \hat{\mathbf{b}}) \cdot (\sqrt{C} \mathbf{u}^1)}{\|\mathbf{b}^* - \hat{\mathbf{b}}\| \|\sqrt{C} \mathbf{u}^1\|} = \frac{(\mathbf{b}^* - \hat{\mathbf{b}}) \cdot (\mathbf{u}^1)}{\sqrt{C}}, \quad (11)$$

where the last equality follows because, at the optimum, $\|\mathbf{b}^* - \hat{\mathbf{b}}\|^2 = C$. At the optimal intervention, by Theorem 1,

$$\underline{b}_\ell^* - \hat{b}_\ell = \frac{w\alpha_\ell}{\mu - w\alpha_\ell} \hat{b}_\ell;$$

now, using the definition $\underline{\mathbf{b}} = \mathbf{U}^T \mathbf{b}$, we have that

$$b_i^* - \hat{b}_i = w \sum_{\ell=1}^n u_\ell^i \frac{\alpha_\ell}{\mu - w\alpha_\ell} \hat{b}_\ell$$

and therefore

$$(\mathbf{b}^* - \hat{\mathbf{b}}) \cdot \mathbf{u}^1 = \sum_i \sum_{\ell=1}^n u_i^1 u_\ell^i \frac{w\alpha_\ell}{\mu - w\alpha_\ell} \hat{b}_\ell = \sum_{\ell=1}^n \frac{w\alpha_\ell}{\mu - w\alpha_\ell} \hat{b}_\ell (\mathbf{u}^1 \cdot \mathbf{u}^\ell) = \frac{w\alpha_1}{\mu - w\alpha_1} \hat{b}_1$$

Hence, using this in equation 11, we can deduce that

$$\rho(\Delta \mathbf{b}^*, \mathbf{u}^1) = \frac{1}{\sqrt{C}} \frac{w\alpha_1}{\mu - w\alpha_1} \hat{b}_1 \geq \sqrt{1 - \epsilon} \quad \text{iff} \quad \left(\frac{w\alpha_1}{\mu - w\alpha_1} \right)^2 \hat{b}_1^2 - C(1 - \epsilon) \geq 0. \quad (12)$$

The following lemma shows that the inequality after the ‘‘if and only if’’ follows from our hypothesis that

$$C > \frac{2\|\hat{\mathbf{b}}\|^2}{\epsilon} \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2,$$

and thus establishing it completes the proof.

Lemma 1. Assume

$$C > \frac{2\|\hat{\mathbf{b}}\|^2}{\epsilon} \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2.$$

Then

$$\left(\frac{w\alpha_1}{\mu - w\alpha_1} \right)^2 \hat{b}_1^2 \geq C(1 - \epsilon) \quad (13)$$

Proof of Lemma 1. Note that

$$C > \frac{2\|\hat{\mathbf{b}}\|^2}{\epsilon} \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2 \implies \epsilon C > \|\hat{\mathbf{b}}\|^2 \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2,$$

and therefore

$$C(1 - \epsilon) < C - \|\hat{\mathbf{b}}\|^2 \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2. \quad (14)$$

But then we have the following chain of statements, explained immediately after the display:

$$\begin{aligned} \left(\frac{w\alpha_1}{\mu - w\alpha_1} \right)^2 \hat{b}_1^2 - C(1 - \epsilon) &\geq \left(\frac{w\alpha_1}{\mu - w\alpha_1} \right)^2 \hat{b}_1^2 - C + \|\hat{\mathbf{b}}\|^2 \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2 \\ &= \left(\frac{w\alpha_1}{\mu - w\alpha_1} \right)^2 \hat{b}_1^2 - \sum_{\ell=1}^n \left(\frac{w\alpha_\ell}{\mu - w\alpha_\ell} \right)^2 \hat{b}_\ell^2 + \|\hat{\mathbf{b}}\|^2 \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2 \\ &= \|\hat{\mathbf{b}}\|^2 \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2 - \sum_{\ell=2}^n \left(\frac{w\alpha_\ell}{\mu - w\alpha_\ell} \right)^2 \hat{b}_\ell^2 \\ &= \left(\frac{\alpha_2}{\alpha_1 - \alpha_2} \right)^2 \sum_{\ell=1}^n \hat{b}_\ell^2 - \sum_{\ell=2}^n \left(\frac{w\alpha_\ell}{\mu - w\alpha_\ell} \right)^2 \hat{b}_\ell^2 > 0. \end{aligned}$$

The first inequality follows from substituting the upper bound on $C(1 - \epsilon)$, statement (14) above, which we derived from our initial condition on C . The equality after that follows by substituting the condition on the binding budget constraint at the optimum, which we derived in Theorem 1. The next equality follows by isolating the term for the first component in the sum and by noticing that that cancels with the first term. The next equality follows by noticing that $\|\hat{\mathbf{b}}\|^2 = \|\underline{\hat{\mathbf{b}}}\|^2$. The final inequality follows because, from the facts that $\mu > w\alpha_1$ and that $\alpha_1 > \alpha_2 > \dots > \alpha_n$, we can deduce that for each $\ell > 1$

$$\frac{w\alpha_\ell}{\mu - w\alpha_\ell} < \frac{w\alpha_\ell}{w\alpha_1 - w\alpha_\ell} = \frac{\alpha_\ell}{\alpha_1 - \alpha_\ell} < \frac{\alpha_2}{\alpha_1 - \alpha_2} \quad \square$$

This concludes the proof of Proposition 2. □