The Leverage of Weak Ties

How Linking Groups Affects Inequality*

Benjamin Golub† Carlos Lever‡

December 2010

Abstract

Centrality measures based on eigenvectors are important in models of how networks affect investment decisions, the transmission of information, and the provision of local public goods. We fully characterize how the centrality of each member of a society changes when initially disconnected groups begin interacting with each other via a new bridging link. Arbitrarily weak intergroup connections can have arbitrarily large effects on the distribution of centrality. For instance, if a high-centrality member of one group begins interacting symmetrically with a low-centrality member of another, the latter group has the larger centrality in the combined network in *inverse* proportion to the centrality of its emissary! We also find that agents who form the intergroup link, the "bridge agents", become relatively more central within their own groups, while other intragroup centrality ratios remain unchanged.

Keywords: social networks, peer effects, centrality, homophily, inequality, weak ties, structural holes, Bonacich centrality.

JEL Classification Numbers: J24, D62, D63

*We thank Matt Elliott, Matt Jackson, Yair Levine, Tomas Rodriguez-Barraquer, Juuso Toikka and seminar participants at Stanford for helpful conversations.

†Graduate School of Business, Stanford University. Email: bgolub@stanford.edu, http://www.stanford.edu/∼bgolub/.

‡Department of Economics, Stanford University.
Email: carloslever@stanford.edu, http://www.stanford.edu/∼cleverg/

1 Introduction

Why do some groups in society invest so much in certain activities — such as education or crime — compared with other groups? How does the distribution of investment depend on social structure? In particular, how does it depend on the nature of the interactions between mostly separate, weakly interlinked groups in society?

Several studies — most prominently Ballester, Calvó-Armengol, and Zenou (2006) and Calvó-Armengol, Patacchini, and Zenou (2009) — have found that network centrality relates to investment levels in social games with linear externalities over a network structure, both in theory and in the data. Thus, we address the above questions and others by studying how network centrality is affected when previously separate groups are first connected.

Our results show that, within the Calvó-Armengol, Patacchini, and Zenou model of investment choices in the presence of social network complementarities, asymmetries in the weak ties that link homophilous groups strongly affect group-level asymmetries in investment. The details of the slight intergroup connections matter in subtle and extremely stark ways for the distribution of investment; the weak ties have a lot of leverage. In particular, even when the intrinsic productive capacities of all individuals are identical, the shares of investment undertaken by different equally-sized groups can be arbitrarily asymmetric as a result of arbitrarily weak intergroup interactions.

We identify two main effects. Suppose we begin with two previously non-interacting groups X and Y and create an intergroup connection between a member x of group X and a member y of group Y. If x was making a small investment relative to his own group and y was making a large one relative to her group, then group X will, as a result, have a *larger* share of the overall equilibrium investment in the newly connected society. Indeed, the ratio of the overall effort level of group X to that of Y will be approximately proportional to the ratio of the pre-connection effort level of y to to that of x. A second, different effect is that if x benefits more from the connection to y than vice versa, then group X will, as a result, have a proportionally larger share of the overall investment. Both effects remain strong no matter how weak the link between groups is; only the asymmetry in the relationship matters.

The two effects complement each other, leading to even greater inequality when both types of asymmetry are present; we fully characterize the impact of both phenomena on investment shares with a simple formula. We also fully characterize how each individual's investment changes and find that the agents with the intergroup connections increase their shares of investment relative to their groups, while the ratios of investments of other agents within a given group do not change.

There are two ways to interpret our results: as a warning about how inequality can depend crucially on the structure of social relationships on inequality across sparsely connected groups; or as a warning against our modeling approach for sparsely connected groups.

From the modeling perspective, network models with strong linear spillovers have become very common in the study of social networks.[1] These models are popular because they give a tractable framework with very little assumptions on the size or structure of the network. Currently the biggest challenge of the theoretical models of economic networks is finding tractable models that allow for complicated patterns of influence across large networks.

Our results give a dire warning for any attempt to bring these models to the data. The relative centrality across groups that are sparsely connected is very sensitive to measurement error in the strength of the links across groups. Unless the researcher is very confident of these measures, the measured centralities must be treated with extreme care.

Our finding isn't completely negative. Our results also show that the relative centrality inside groups that are well connected is robust to perturbations that connect the group to other parts of the network and to perturbations in the link structure inside the group.

From the inequality perspective, an important characteristic we observe in many social networks is that people tend to form relationships with people who are similar to them by race, gender, age, class, geographical location and other factors.[2] This effect is called homophily in the sociological literature and while it is widely believed that it matters for social outcomes, there is not much of a theoretical understanding of exactly *how* it matters.

Our results suggests that the details of the links bridging homophilous groups can be very important to determine relative outcomes of groups in societies. These details include the identity of the individuals that have relationships across groups and the position they have in the network of their own group, as well as the relative direct influence between bridge agents of different groups.

Inequality across different social groups in society has been widely documented. For instance, in data from the National Longitudinal Survey of Youth (1971–91 waves), 52 percent of white American males have completed high school at age 24, but only 38 percent of Latin Americans males have done so.[3] As such, we cannot rule out that the structure of social relationships is influencing the outcomes.

On the normative side, there are many social interventions and programs whose explicit

[1] Ballester et al. (2006); Calvó-Armengol et al. (2009); DeMarzo et al. (2003); Bramoullé et al. (2009); DeGroot (1974); Golub and Jackson (2010, 2009); Lever (2010).

[2] See McPherson et al. (2001) for a comprehensive survey.

[3] See (Cameron and Heckman, 2001) for this and more statistics on racial gaps in school achievement and an overview of some of the literature.

purpose is to link separate social groups. This includes cultural exchange between countries as well as charities such as the Big Brothers and Big Sisters of America, which pairs vulnerable youth with mentors from different social groups. In an equilibrium model, such interventions have indirect effects on many agents beyond the ones they directly involve, and ultimately on inequality of investment within a society. Understanding these externalities is important to designing such of interventions, and our work takes a step toward such an understanding.

On the way to our results, we make two general contributions that are of independent interest for the study of centrality measures in networks.

First, we show that eigenvector centrality is the limit of Bonacich centrality in an appropriate sense. This extends a result proved in Bonacich (1991) to a much larger class of interaction matrices. It also establishes a broad connection between two of the most important centrality measures used in practice: Bonacich centrality, on the one hand, and eigenvector centrality, on the other.[4]

Second, we derive a characterization of the eigenvector centrality of a network obtained by connecting two previously disconnected networks. Beyond the application to a model of investment with complementarities, this characterization applies to eigenvector centralities in other settings such as the Google PageRank measure, the theory of Markov chains (where the eigenvector centrality vector is the stationary distribution) and numerous applications in sociology. In these settings, too, arbitrarily weak intergroup links can have arbitrarily large effects on the distribution of centrality across groups. We comment on some of the broader implications of this in Section 6.

The intergroup connections that we study have been a focus of investigation across several fields. Granovetter's (1973) seminal analysis of how social networks mediate learning about jobs emphasized the informational importance of weak ties connecting otherwise separate communities. Burt (1992) developed an influential theory of structural holes, exploring the advantaged positions enjoyed by the agents incident to a bridge — a feature that is also present, for different reasons, in our analysis. Both works have spawned large literatures in sociology. In a recent paper on the spread of misinformation in social networks by Acemoglu, Ozdaglar, and ParandehGheibi (2009), bridges are important in characterizing when agents can successfully spread their potentially biased beliefs. Bridges in graphs are also well-studied objects in graph theory (see, for example Diestel (2005)). Our work links the notion of bridges with the study of centrality which, as mentioned above, has been an important

[4]Bonacich centrality was introduced in Bonachich (1987). Eigenvector centrality is also known as Katz prestige (Katz, 1953). It measures influence in an important model of belief evolution (French, 1956; DeGroot, 1974), and is also used in the computation of the Google PageRank measure (Langville and Meyer, 2006).

area in sociology and network analysis more generally (Wasserman and Faust, 1994).

The paper is organized as follows. Section 2 sets up the model, presents the main definitions on networks and provides our main result: a formula for how connecting groups changes eigenvector centrality. Section 3 then uses the result to derive new predictions for two economic applications, investment decisions with positive spillovers through a social network (Section 3.1) and consumption decisions with social influence (Section 3.2). In Section 6, we summarize the main insights and comment on some implications and extensions of the results.

2 How Linking Groups Affects Centralities

2.1 Setup

Let there be n agents (also called nodes) $N = \{1, \ldots, n\}$. A weighted social network is a matrix $\mathbf{A}$ with non-negative entries a_{ij}. An entry a_{ij} represents the weight of the link from i to j. We will assume that $\mathbf{A}$ is a column-stochastic matrix, $\sum_i a_{ij} = 1$ for all j. (See Section 3.1 and Section 3.2 for an economic justification.) We assume that $a_{ii} = 0$ for all i, which means we don't allow self-links.

Below are some common definitions for networks.[5]

A *walk* in $\mathbf{A}$ is a sequence of nodes $i_1, i_2, \ldots, i_K$, not necessarily distinct, such that $a_{i_k i_{k+1}} > 0$ for each $k \in \{1, \ldots, K-1\}$. We say that a walk $i_1, i_2, \ldots, i_K$ goes *from* i_1 *to* i_K. The *length* of the walk is defined to be $K - 1$. The *weight* of the walk is defined to be the product $\prod_{k=1}^{K-1} a_{i_k i_{k+1}}$. A walk is *closed* if it starts and ends at the same node, i.e. $i_1 = i_K$.

A *path* is a walk whose nodes are distinct. A network $\mathbf{A}$ is *path-connected* (also called *irreducible*) if for every pair of agents i, j, there is a directed path from i to j and back.

We are interested in studying changes in eigenvector centrality, as defined below.

Definition 1 (Eigenvector centrality). Given a nonnegative, path-connected matrix $\mathbf{A}$, the eigenvector centrality $\mathbf{e}(\mathbf{A})$ is the right-hand (column) eigenvector of $\mathbf{A}$ with nonnegative entries. This eigenvector is unique up to scale; we will always assume its entries sum to 1.[6]

We may also speak of the eigenvector centrality of a group $X \subset N$, written $e_X(\mathbf{A})$, which is just the sum of the eigenvector centralities of its members.

[5] See Jackson (2008) for these and other definitions relating to networks.

[6] All vectors in the paper are column vectors unless otherwise noted. The existence and uniqueness of the eigenvector centrality are a consequence of the Perron-Frobenius theory of nonnegative matrices, and are discussed in (Meyer, 2000, Section 8.3).

In many economic applications, we are interested in studying how the centrality of an agent changes when we change the weight of a link or when we connect disconnected groups. When a network is path-connected, there are formulas for the derivative of the eigenvector centralities with respect to the weight of any link; we discuss some of these in Section 5. Unfortunately these formulas do not apply for disconnected networks. Since linking disconnected groups features prominently in studies of trade, peer effects and strategic network formation, it's important to have tools to address this scenario.

Our main result below bridges this gap and provides a closed-form solution for the change in eigenvector centrality when we link disconnected groups. To get closed-form solutions we focus on a single link connecting two segments of society. This *two-way bridge* is defined formally below.

Definition 2 (A two-way bridge). Let X, Y be a partition of the agents. We say a path-connected matrix $\mathbf{A}$ has a *two-way bridge* if there exist $x \in X$ and $y \in Y$ so that

$$a_{xy} > 0, a_{yx} > 0$$

and

$$a_{ij} = 0; \forall (i,j) \in (X \times Y) \setminus \{(x,y)\}$$

We call x and y the *bridge nodes.*

In words, $\mathbf{A}$ has a two-way bridge if it can be divided into two path-connected groups X and Y such that there is a single pair of agents (x, y) who have influence across groups. Having a single link that connects the groups will be crucial to get a clean closed-form solution, but we believe most of our qualitative results hold if the groups have many bridge links but are sparsely connected. We discuss more on these matters in Section 5.

It will be important for our purposes to talk about the restriction of the network to a subgroup. Given a matrix $\mathbf{A}$ with a two-way bridge between X and Y and bridge nodes $x \in X$ and $y \in Y$, let $\mathbf{A}^X$ denote the matrix with index set X so that

$$a_{ij}^X = \begin{cases} a_{ij} & \text{if } j \neq x \\ a_{ij}/(1 - a_{yx}) & \text{if } j = x \end{cases}.$$

The entries are normalized to make $\mathbf{A}^X$ column-stochastic while preserving the relative influence of x over the other members of X. We define $\mathbf{A}^Y$ analogously. These restrictions of $\mathbf{A}$ capture the interactions that would exist if there were no bridge between X and Y.

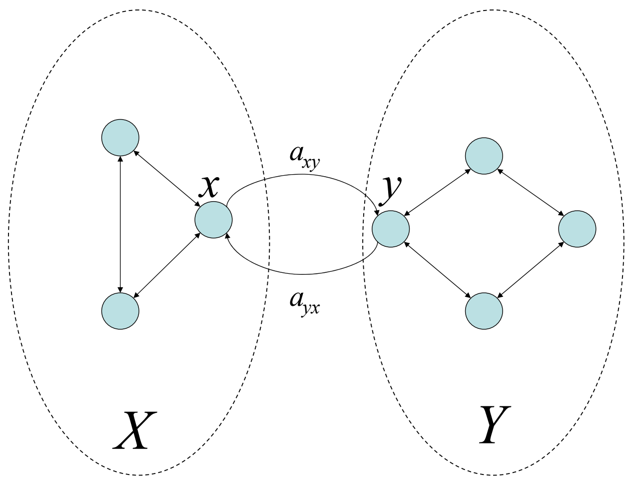

Figure 1: The nodes (x, y) form a two-way bridge between groups X and Y.

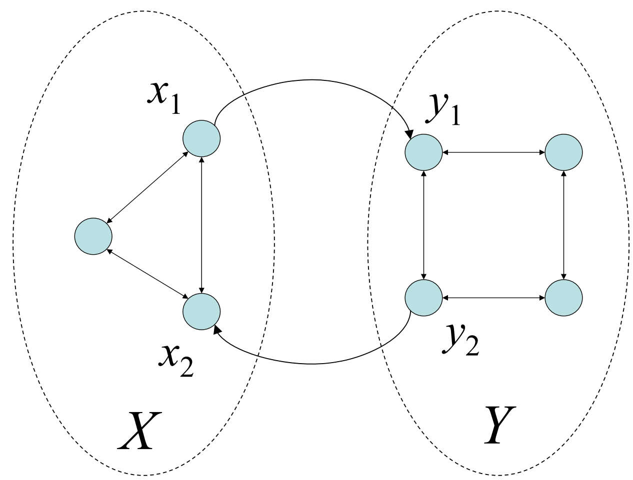

Figure 2: This is not a two-way bridge, because the node that connects X to Y is not the same that connects Y to X.

Intuitively, we can think that in the beginning X and Y belong to separate networks each with column-stochastic interaction matrices $\mathbf{A}^X$ and $\mathbf{A}^Y$. Then we connect them by creating the link between x and y, giving rise to a new combined interaction matrix $\mathbf{A}$.

2.2 The Main Result

We now present our main result. The following theorem characterizes the eigenvector centralities of all agents in the connected network based on their eigenvector centralities in the restricted networks.

THEOREM **1.** *Assume* $\mathbf{A}$ *is path-connected and has a two-way bridge between* X *and* Y *with bridge nodes* $x \in X$ *and* $y \in Y$*. Let*

$$c_x = 1 - a_{yx}(1 - e_x(\mathbf{A}^X)) \text{ and } c_y = 1 - a_{xy}(1 - e_y(\mathbf{A}^Y))$$

and

$$r = \frac{a_{xy}}{a_{yx}} \cdot \frac{e_y(\mathbf{A}^Y)}{e_x(\mathbf{A}^X)} \cdot \frac{c_x}{c_y}. \tag{1}$$

Then

$$e_i(\mathbf{A}) = \frac{1}{1+r} \cdot \begin{cases} e_i(\mathbf{A}^X) \cdot r c_x^{-1} & \textit{if } i = x \\ e_i(\mathbf{A}^X) \cdot r c_x^{-1}(1 - a_{yx}) & \textit{if } i \in X \setminus \{x\} \\ e_i(\mathbf{A}^Y) \cdot c_y^{-1} & \textit{if } i = y \\ e_i(\mathbf{A}^Y) \cdot c_y^{-1}(1 - a_{xy}) & \textit{if } i \in Y \setminus \{y\}. \end{cases}$$

The formulas are simple to interpret. The centrality of an agent in the connected network is proportional to the his centrality in the disconnected network multiplied by certain correction factors that depend on the weights of the two-way bridge. The agents in in X have their centralities multiplied by a factor of r, and then adjusted by the correction factor c_x^{-1}. The agents in in Y have an analogous correction term. The fraction $(1 + r)^{-1}$ on the outside is merely a normalization.

Notice that the ratios of centralities between members of the same group remain the same except when a bridge node is involved. Relative to the other agents in his group, the bridge agent x has his centrality increased by a factor of $(1 - a_{yx})^{-1}$. This generates a perverse incentive: when agent x increases the benefits he gives to y in the other group, all members of his group lose centrality, but x gains a higher relative position inside of X.

The logic for the proof is straightforward. Since we know the eigenvector is unique we guess it has the form above and verify it. The real challenge was coming up with the candidate guess. We present the intuition below. Reviewing this will clarify why our approach requires connecting X and Y through a single two-way bridge.

Extending our approach to multiple links is difficult, but once the groups have become connected by single-link, we can apply derivative formulas for path-connected networks. See Section 5.

2.3 The Intuition for the Formula

The key to guess the formula for eigenvector centrality is to use the fact that the eigenvector centrality of $\mathbf{A}$ is the same as the stationary distribution of a Markov chain with a transition matrix which is the transpose of $\mathbf{A}$. This Markov chain is one in which a particle randomly hops around the nodes N with a probability a_{ji} of moving from node i to node j. While the intuitions of this process are rather far from the economic applications we present below, the advantage of the Markov chain approach is that we can use probabilistic results to study our problem. This provides natural intuitions that motivate the expression.

Suppose at time $n = 0$ the particle starts a random walk from node x. Let W_n be the position of the particle at time n, which is a random variable. Analogously define a Markov chain starting at x corresponding to $\mathbf{A}^X$ and denote the position of that particle at time n in that process by W_n^X. This process corresponds to a particle that walks only in X, with the transition probabilities prescribed by $\mathbf{A}^X$.

Define the random variable

$$T_i = \inf\{n \geq 1 : W_n = i\}$$

This is the time of the first visit to i after time 0. A well-known formula (Durrett, 2005, Chapter 5, (4.3)) states that for any $M \subseteq N$

$$e_M(\mathbf{A}) = \frac{\mathbb{E}_x\left[\sum_{n=0}^{T_x-1} 1_{\{W_n \in M\}}\right]}{\mathbb{E}_x(T_x)}. \tag{2}$$

Here, the subscripts on the expectations remind us that the particle starts at x and $1_{\{W_n \in M\}}$ is the indicator random variable which takes value 1 if and only if the particle is in M at time n.

This tells us, for example, that:

Claim 1. *For $i \in X$ such that $i \neq x$*

$$\frac{e_i(\mathbf{A})}{e_x(\mathbf{A})} = (1 - a_{yx}) \cdot \frac{e_i(\mathbf{A}^X)}{e_x(\mathbf{A}^X)}.$$

This is because (2) applied to $M = \{x\}$ yields:

$$e_x(\mathbf{A}) = \frac{1}{\mathbb{E}_x(T_x)}. \tag{3}$$

When the formula is applied to $M = \{i\}$, we reason as follows. After leaving x, the particle does one of two things. (1) It may go to y with probability a_{yx}, in which case it cannot possibly hit i before its first return to x (recall that there is a bridge from x to y which, by definition, is the only way to get between the two groups). (2) Or it may go somewhere within X with probability $(1 - a_{yx})$, and conditional on this event its travels throughout X before its next return to x have the same probability distribution as they did without the bridge. Thus we may write

$$e_i(\mathbf{A}) = \frac{(1 - a_{yx})\mathbb{E}_x^X \left[\sum_{n=0}^{T_x - 1} 1_{\{W_n^X = i\}}\right]}{\mathbb{E}_x(T_x)}. \tag{4}$$

The superscript on the expectation indicates that it refers to the chain conditional on staying in X, whose transition matrix is given by the transpose of $\mathbf{A}^X$. Dividing (4) by (3) and using (2) for the chain corresponding to $\mathbf{A}^X$ yields the claim. An analogous claim holds for Y.

This pins down the new eigenvector centralities inside each group. To finish we only have to determine the relative influence across groups. For this it suffices to consider the eigenvector condition

$$e_x(\mathbf{A}) = \sum_j a_{xj} e_j(\mathbf{A}).$$

3 Motivating Examples

We know use our formula to provide new results for two economic problems. In both examples, eigenvector centrality will characterize the equilibrium decisions in the limit with strong social influence.

3.1 Investment Decisions with Strong Social Spillovers

We work with a special case of the model of Ballester, Calvó-Armengol, and Zenou (2006) and define the following investment game. Each agents i selects an effort level $z_i \geq 0$, which can be interpreted as human capital investment. Agent i selects z_i to maximize:

$$u_i(z_1, \cdots, z_n) = z_i - \frac{1}{2} c z_i^2 + \gamma \sum_{j \neq i} a_{ij} z_i z_j. \tag{5}$$

Here, $c > 0$ is a marginal cost parameter; $\gamma > 0$ is a social spillover parameter and a_{ij} is an entry of the column-stochastic matrix $\mathbf{A}$.

The interpretation is that each person's utility is a combination of three things. First, there is a linear own-effort effect which we normalize to have a unity coefficient. Second, there is a convex cost in own effort introduced by the quadratic second term and parametrized by c. We assume everybody has the same marginal cost of effort. Finally, there are social complementarities. An agent j spends an amount of time a_{ij} "teaching" agent i. The teaching has complementarities in the level of knowledge of the teacher and of the student. The benefit i receives from j is increasing in z_i and in z_j. The total benefit i receives from other people's "lessons" is given by $\gamma \sum_j a_{ij} z_i z_j$. Agent i's marginal benefit of investing in education is increasing in the investment level of the people connected to her through the social network.

We assume that each agent j spends the same amount of time teaching other agents.[7] With an appropriate choice of γ we normalize this so that $\sum_j a_{ij} = 1$ and the matrix $\mathbf{A}$ is column-stochastic. This amounts to assuming that no agent has more potential to teach than others, and corresponds to our focus on the network creating inequality as opposed to ex-ante differences between agents. Note, however, that this is not a strong symmetry assumption. The weight from i to j might be different from the weight from j to i. The returning link might even not exist at all.

In order to focus on the consequences of changing the network, we will treat the values for a_{ij} as exogenous.

Theorem 2 below characterizes the equilibria of the game. It is almost follows from Theorem 1 in Ballester, Calvó-Armengol, and Zenou (2006). Before the statement, we need one definition.

Definition **3.** Given a nonnegative matrix $\mathbf{M}$, provided that $[\mathbf{I} - \alpha\mathbf{M}]^{-1}$ is well-defined and nonnegative, the vector of Bonacich centralities of $\mathbf{M}$ with parameter α is

$$\mathbf{b}(\mathbf{M}, \alpha) = [\mathbf{I} - a\mathbf{M}]^{-1}\mathbf{1} = \sum_{k=0}^{\infty} \alpha^k \mathbf{M}^k \mathbf{1}, \tag{6}$$

where $\mathbf{1}$ is the $n \times 1$ vector of ones. Let $B(\mathbf{M}, \alpha)$ be the sum of entries of $\mathbf{b}(\mathbf{M}, \alpha)$.

Theorem **2.** *The investment game has a Nash equilibrium if and only if $\frac{\gamma}{c} < 1$. When it exists, the Nash equilibrium is unique and is given by:*

$$\mathbf{z}^* = \frac{1}{c}\mathbf{b}\left(\mathbf{A}, \frac{\gamma}{c}\right).$$

All proofs appear in the Appendix.

The intuition for this result is that investment is proportional to network position as measured by Bonacich centrality. The agents who invest the most are the ones who most benefit from feedback loops of social complementarities. This can be seen in equation (6) which shows that the Bonacich centrality of agent i is a weighted sum of the walks in $\mathbf{A}$ that start at i. This is just a multiplier effect. The size of the multiplier is determined by the

[7] A model where every agent spends the same fixed-amount to "learn" rather than "teach" would be strategically equivalent to out second motivating example. See Section 3.2.

structure of the network and each agent's position determines how much he or she benefits.[8]

Bonacich and eigenvector centrality are closely related: as the network feedback-loops become large, Bonacich centrality converges to eigenvector centrality. This is stated generally in Theorem 3, which extends a theorem in Bonacich (1991). Bonacich shows the result for symmetric matrices using a diagonalization argument. Our proof is quite different and does not require symmetry or even diagonalizability.

To state the result precisely and in full generality, we need one more technical notion — that of *aperiodicity*, which is standard in the study of Markov chains. It is a mild technical condition that holds generically and requires that the greatest common divisor of the lengths of all closed walks be 1. For example, this condition holds if there is at least one closed walk with two agents and at least one closed walk with three agents.

Theorem **3.** *Given a nonnegative path-connected aperiodic matrix* $\mathbf{M}$ *whose largest eigenvalue in magnitude is* μ*, we have*

$$\lim_{\alpha \uparrow \mu^{-1}} \frac{\mathbf{b}(\mathbf{M}, \alpha)}{B(\mathbf{M}, \alpha)} = \mathbf{e}(\mathbf{M}).$$

The relationship between Bonacich centrality and eigenvector centrality can be understood in the following way: the Bonacich centrality of i is computed by starting with a baseline centrality of 1 (which corresponds to the linear own-effort term in our investment game) and sums all walks starting at i, with walks of length k getting weight α^k.

Eigenvector centrality measures influence by giving equal weight to all walks starting at i. The number of such paths is infinite, but by taking appropriate limits, we can still make comparisons between nodes. Eigenvector centrality pins down these centralities. Adding a normalization determines the levels.

The higher α is, the greater is the importance of long walks for Bonacich centrality. In the limit, the baseline effect and the short-distance walks are completely insignificant. Therefore as the network feedback become large, the ratio of Bonacich centralities converges to the ratio of eigenvector centralities.

This allows us to use our result on eigenvector centrality to study the investment game when there are strong social spillovers, focusing on the distribution of investment across society. For a given interaction matrix $\mathbf{A}$ and parameters γ, c, let the equilibrium *investment*

[8]When $\frac{\gamma}{c} \geqslant 1$, the social spillovers are so big that there is no equilibrium, because agents would always want to invest more. One way to see this is to note that this is a supermodular game, so the best-response mapping converges to the lowest Nash equilibrium when the mapping starts from the lowest action. When $\frac{\gamma}{c} \geqslant 1$, this dynamic is explosive, so no equilibrium can exist.

share of a set X of agents be defined by

$$s_X(\mathbf{A}, \gamma/c) = \frac{\sum_{i\in X} z_i^*}{\sum_{i\in N} z_i^*},$$

where $\mathbf{z}^*$ is the equilibrium as characterized in Theorem 2. Note that even though $\mathbf{z}^*$ itself depends on the levels of both γ and c, the equilibrium shares as defined above depend only on the ratio γ/c, which explains the notation. We will also write $s_i(\mathbf{A}, \gamma/c)$ instead of $s_{\{i\}}(\mathbf{A}, \gamma/c)$, and use the boldface notation $\mathbf{s}(\mathbf{A}, \gamma/c)$ for the vector with $s_i(\mathbf{A}, \gamma/c)$ in the ith position. We will drop the arguments when they are clear from context.

A corollary of Theorem 3 is that when social spillovers are big, these shares are well approximated by eigenvector centrality.

Corollary **1.** *Assume* $\mathbf{A}$ *is path-connected and aperiodic. As* γ/c *approaches* 1 *from below, the investment share of a group* X *approaches the eigenvector centrality of* X *with respect to* $\mathbf{A}$. *That is:*

$$\lim_{\frac{\gamma}{c}\uparrow 1} \mathbf{s}(\mathbf{A}, \gamma/c) = \mathbf{e}(\mathbf{A}).$$

As a consequence of Theorem 1, we can characterize the relative investment shares of the two groups.

Corollary **2.** *Assume* $\mathbf{A}$ *is path-connected and has a two-way bridge between* X *and* Y *with bridge nodes* $x \in X$ *and* $y \in Y$. *Then*

$$\frac{e_X(\mathbf{A})}{e_Y(\mathbf{A})} = \frac{a_{xy}}{a_{yx}} \cdot \frac{e_y(\mathbf{A}^Y)}{e_x(\mathbf{A}^X)} \cdot \frac{1 - a_{yx}(1 - e_x(\mathbf{A}^X))}{1 - a_{xy}(1 - e_y(\mathbf{A}^Y))}.$$

In view of Corollary 1, this means that the ratio of investment shares is given by the above equation. A simpler formula can be obtained by taking a limit as the link between X and Y becomes weak.

Corollary **3.** *Let* $\mathbf{A}$ *be as above and* $a_{xy} = k a_{yx}$ *for some constant* $k > 0$. *Then*

$$\lim_{a_{yx}\to 0} \frac{e_X(\mathbf{A})}{e_Y(\mathbf{A})} = k \cdot \frac{e_y(\mathbf{A}^Y)}{e_x(\mathbf{A}^X)}.$$

This simplified formula contains the main insights about how the bridge affects each group's centrality.

First, the ratio of the investment of group X to the investment of group Y is directly proportional to the ratio of the weights between the bridge-nodes: $({}^{a_{xy}}/_{a_{yx}})$. If agent x gets more benefit from the complementarity with y than vice versa, then group X ends up doing most of the investment in the society after the connection is made.

Second, the ratio of the investment of group X to the investment of group Y is directly proportional to the ratio ${e_y(\mathbf{A}^Y)}/{e_x(\mathbf{A}^X)}$. This is the ratio of the bridge agents' investment shares in their own groups before the bridge is created. In other words, connecting a relatively high-investing member of group Y to a relatively low-investing member of group X is good for the investment share of group X in the combined network.

Third, and perhaps most surprisingly, arbitrarily weak links can have arbitrarily large effects on the investment shares in the combined network. Even in the limit as the level of interaction between the two groups tends to zero, the shares of investment can remain quite unbalanced. Indeed, the formula in Corollary 3 shows that there is a discontinuity as a link is added; without the link, two groups may invest the same amounts, but after an arbitrarily weak link is introduced, they may become very unequal.

Fourth, the investment ratio of the groups X and Y depends only on the link between them and on the within-group investment shares of the bridge nodes. The relative size of each group is irrelevant as well as other traditional metrics on network. It does not matter which group is more densely connected or which group has a larger network diameter.

This inequality in investment does not go away if the link between both groups is strong. The next theorem shows that the comparative statics for strong links go in the same direction: the investment share of X relative to Y increases if y increases the amount of help she gives to x; it is also increasing in the influence of y within his own group Y.

THEOREM **4.** *The ratio of investments* $\frac{e_X(\mathbf{A})}{e_Y(\mathbf{A})}$ *is strictly increasing in* a_{xy} *and* $e_y(\mathbf{A}^Y)$*. It is strictly decreasing in* a_{yx} *and* $e_x(\mathbf{A}^X)$*.*

3.2 Consumption with Strong Social Influence

We now study a model of consumption with social influence. Every agent makes a consumption decision $z_i \in \mathbb{R}$ and is influenced by the choice of his neighbors. Agents simultaneously make their consumption decision to maximize:

$$u_i(z_i, \mathbf{z}_{-i}) = -\frac{1}{2}\big(z_i - \theta_i\big)^2 - \frac{\beta}{2}\big(z_i - \sum_j a_{ij} z_j\big)^2$$

The parameter θ_i represents agent i's ideal choice in a world without social influence. It is fixed and common knowledge. We call it the autarky ideal point of i. The weight β represents the agent's preference for acting like his neighbors. Each parameter a_{ij} is an entry of a matrix of influence $\mathbf{A}$ and represents the influence of j over i. We assume the matrix of influence A is row-stochastic, so every agent is trying to match a weighted average of the decisions of his neighbors.

Agents face a trade-off between being closer to their autarkic ideal point and being close to their neighbors' decision. Because being away from either point has an the increasing marginal cost, agents chose something in-between.

The theorem below characterizes the equilibria of the game:

THEOREM **5.** *The unique Nash equilibrium of the game is*

$$\mathbf{z}^* = \frac{1}{1+\beta}\left[\mathbf{I} - \frac{\beta}{1+\beta}\mathbf{M}\right]^{-1}(\theta_1, \ldots, \theta_N)'.$$

COROLLARY **4.** *In the limit when social preferences dominate ($\beta \to \infty$), all agents make the same consumption choice which is a weighted-sum of the autarkic ideal points:*

$$\lim_{\beta\to\infty} z_i^* = \sum_j e_j(\mathbf{A})\theta_j; \forall i.$$

The equilibrium of the consumption game is very similar to that of the investment game. The strategies in Theorem 5 involve a modified version Bonacich centrality. There is one important difference between the games: in the investment game with strong social spillovers, each agent invests different amounts in proportion his eigenvector centrality. In the consumption game with strong social spillovers, all agents invest the same amount, which is a weighted sum of the autarkic ideal points. The weights are the eigenvector centralities.

Using Theorem 2, we see that when group X and Y start interacting (become linked) the new equilibrium can be very biased toward the ideal-points of one of the groups. In particular if $a_{yx} \geqslant a_{xy}$ and $e_y(A^Y) \geqslant e_x(A^X)$ with one strict inequality, the ideal-points of group X will have a larger weight in the final consumption decision.[9]

In the limit when the bridge-links are close to zero as in Corollary 3, the new consumption is

$$z^* = \frac{r}{1+r}\sum_{i\in X} e_i(A^X)\theta_i + \frac{1}{1+r}\sum_{j\in Y} e_j(A^Y)\theta_j$$

where r is the ratio

$$r = \frac{a_{yx}e_y(A^Y)}{a_{xy}e_x(A^X)}$$

[9] Here and throughout the section a_{yx} and a_{xy} are switched with respect to the formula in Corollary 2 because here the matrix A is row-stochastic, so we have to transpose it to apply Theorem 1.

We can take the analysis one step further by assuming that consumption preferences evolve over time. Suppose that the consumption game is played repeatedly many times, but that the ideal points θ_i change randomly every period. After the θ_i's are realized they become common-knowledge and all agents play the static Nash equilibrium.

Suppose all agents in the same group draw θ_i from the same distribution but that the realizations of θ_i are independent across agents. Let σ_x^2, σ_y^2 be the variances of the distribution for groups X and Y. Suppose that half the members of society belong to group X and half to group Y. Before the groups are connected the average of agents' period-to-period variance in consumption is

$$\sigma_{z^*}^2 = 2\Big[\Big(\frac{1}{2}\Big)^2 \sigma_x^2 \sum_{i \in X} \big(e_i(A^X)\big)^2 + \Big(\frac{1}{2}\Big)^2 \sigma_y^2 \sum_{j \in Y} \big(e_j(A^Y)\big)^2\Big]$$

Assume that the variance in the equilibrium-consumption in group X is larger than in group Y.

$$\sigma_x^2 \sum_{i \in X} \big(e_i(A^X)\big)^2 > \sigma_y^2 \sum_{j \in Y} \big(e_j(A^Y)\big)^2$$

This can happen for two (non-exclusive) reasons: the variance of tastes per agent in X is larger than in Y: $\sigma_x^2 > \sigma_y^2$; or the eigenvector centralities in Y group are more evenly balanced than in X: $\sum_{j \in Y} \big(e_j(A^Y)\big)^2 < \sum_{i \in X} \big(e_i(A^X)\big)^2$. This second effect occurs because social influence between members of the same network "averages out" the shocks in θ_i. But if the eigenvector centrality is distributed very unevenly there is little "averaging-out". In the extreme case, all agents listen to a single member of the group and the group variance in consumption is identical to the variance of that individual's θ_i.

Now assume groups X and Y become connected. The new the variance in consumption becomes

$$\sigma_{z^*}^2 = \sigma_x^2 \sum_{i \in X} \big(e_i(A)\big)^2 + \sigma_y^2 \sum_{j \in Y} \big(e_j(A)\big)^2,$$

which in the limit as $a_{xy}, a_{yx} \to 0$ becomes

$$\sigma_{z^*}^2 = \Big(\frac{r}{1+r}\Big)^2 \sigma_x^2 \sum_{i \in X} \big(e_i(A^X)\big)^2 + \Big(\frac{1}{1+r}\Big)^2 \sigma_y^2 \sum_{j \in Y} \big(e_j(A^Y)\big)^2$$

Notice that there is a 2 that dropped out when we connect the two groups. This represents the drop in variance because society it now twice as large, so it has more members "averages out" the shocks in θ_i. This potential reduction is diminished depending on how we connect the groups. When $r \to 1$ the variance in consumption is just the variance in group y and there is no "averaging-out" effect across groups, so the period-to-period variance increases.

Therefore connecting the groups makes the consumption decisions more homogenous inside each period and might increase the average variance in consumption from period-to-period. This last effect occurs because the tastes of the volatile group has a disproportionate influence in the equilibrium consumption decisions.

4 Empirical Implications: Sensitivity of Centrality Measures

Our results also point out that the global properties of centrality measures can be very sensitive to small, local perturbations when there are islands bridged by weak ties, as there often are in real social networks. Imbalances in these weak links can lead to large global imbalances of centrality, so that small absolute changes in some network parameters can lead to large absolute changes in the centralities.

This has implications for the empirical analysis of networks.[10] When there is noise in the measurement of the network and the network has an "island-like" structure, our analysis suggests that the relative centrality across groups should be treated with caution, since they may be very sensitive to measurement error in the links across groups.

Suppose, for example, that there is one link between agents x and y bridging two groups, and the estimate of the complementarity a_{xy} that agent x gets from agent y is twice the estimate of a_{yx}; however, in reality, the two are equal. This is quite plausible if both quantities a_{xy} and a_{yx} are small, so that the error is very small in absolute terms but leads to a large error in the estimate of the ratio a_{xy}/a_{yx}. In this case, the measurement error causes all the computed centralities in the group of agent x to be off by a factor of two! On the other hand, Theorem 1 shows that the ratios of centralities within a group will be fairly accurate, because they essentially do not depend on the weak intergroup links.

More generally, this seems to suggest that, in the presence of measurement error, eigenvector centrality is not a robust tool for comparing centralities of nodes located in parts of a network that are only weakly connected to each other. At the same time, it is quite reasonable for comparing centralities of agents in a thick region of the network. Making this speculation more precise could be a fruitful direction for further work.

[10]We thank Matt Jackson for emphasizing the implications of this to us.

5 Several Links Between Groups

Our analysis characterizes how investment and consumptions decisions — or, more generally, network centralities — change when previously separate groups start interacting. A natural question is what happens when further links are added beyond the first, so that we are looking at two islands connected by a few links as opposed to just one bridge. This question is answered in a beautiful paper of Conlisk (1985). Start with a path-connected column-stochastic $\mathbf{A}$ and consider perturbating a single link: $A_{ji} \to A_{ji} + \epsilon$ and $A_{ki} \to A_{ki} - \epsilon$. All other entries of $\mathbf{A}$ are held fixed. This corresponds to agent i increasing his complementarity or interaction with j (the "favored node") at the expense of k (the "disfavored node"). The decrease is necessary to satisfy the assumption that the total complementarity shared out by any agent is fixed, so that increasing the benefit to one neighbor requires decreasing the benefit to another.

Conlisk shows that, under this perturbation, for any ℓ, we have

$$\frac{\partial e_\ell}{\partial \epsilon} = e_i(\mathbf{A}) \cdot (w_{j\ell} - w_{k\ell}),$$

where $w_{j\ell}$ is the mean first passage time from j to ℓ in the Markov chain whose transition matrix is given by the transpose of $\mathbf{A}$. (See our intuition in Section 2.3.) This is a measure of the network distance between j and ℓ. Thus, the centrality of node ℓ changes in proportion to the difference $w_{j\ell} - w_{k\ell}$. If ℓ is closer to j (the favored node) she will gain centrality while if she is closer to the disfavored node k she will lose it. The magnitude of the change is also proportional to the centrality of the node i, the agent doing the redistributing.

The intuitions of the formula cohere with those of our main result. Adding directed links, all else equal, helps those who get more weight (and the agents near them) and hurts those who lose weight (and the agents near them). This effect is stronger when the originating node is more influential.

Indeed, in combination with Conlisk's results, our treatment of bridges completes the picture on the comparative statics of the centralities of normalized network matrices (equivalently, the stationary distributions of Markov chains). For any change one might like to consider, the results working in tandem can completely characterize the effects of the perturbation in an intuitive way — though an exact computation of the quantities would require solving a system of differential equations.[11] Of course, since the relevant centralities are solu-

[11] One change which may at first seem tricky to treat is that of introducing a connection between two separate groups which is not a two-way bridge, as in Figure 2. For example, it may be that there is a directed link from $x \in X$ to $y \in Y$, but the only link in the reverse direction comes back from some $y' \neq y$

tions to a well-studied fixed-point problem, one could always simply compute them explicitly with a program like MATLAB. The virtue of the comparative statics formulas is that they explain exactly what matters and how, which sheds more light on positive and normative economic questions than a computational approach. Indeed, these formulas remove some of the mystery of network centralities. Not only do we know that they satisfy some desirable fixed-point property, but we also know, in a reasonably explicit way, how they change when interesting changes happen in the network.

6 Conclusion

We presented a new theoretical result that provides a simple, closed-form solution for how eigenvector centrality changes when disconnected groups become linked. Previous work had studied only the derivative of eigenvector centralities in connected networks.

The result is particularly useful for comparing group centralities. We find that depending on the way the groups are connected, the influence of one group can dominate the influence of the other. In particular, if a member with a central position in group Y becomes linked to a member with a noncentral position in group X, the resulting eigenvector centrality of group X will be proportionally larger than that of group Y. This occurs regardless of the relative size of the groups and other network metrics: their clustering, the density of their connections, or their diameter.

We also showed how the result is useful through two economic applications. In one, individuals invest in human capital and generate positive externalities by teaching others through a social network. In equilibrium, agents invest in proportion to their eigenvector centrality, which means that one group may end up with a disproportionate share of the total investment after the groups become connected.

In the second example, every period agents locate themselves in a one-dimensional consumption space by balancing their private tastes with their desire to match the decisions of their neighbors. In equilibrium, the group with a higher eigenvector centrality has a disproportionate weight in determining the final location of consumption decisions in society. Furthermore, connecting two groups, somewhat counterintuitively, increases the variance

in Y to $x' \in X$, which may or may not be the same as x. Using techniques similar to those used to prove our main result, one can characterize the post-perturbation centralities exactly. The formulas are not particularly elegant. But one can also characterize them by combining Conlisk's result with ours. One would add a two-way bridge at first (say between x and y), and then use the Conlisk formulas to characterize what happens when one of its directions is gradually replaced with the correct return link from y' to x'. We thank Tomás Rodriguez-Barraquer for pointing out that such a change can be substantively important, for example when academic disciplines are first linked via two separate one-way bridges.

in consumption from one period to the next. This happens because a small group has a disproportionate influence in the final decision.

Our results suggest that network models with linear spillovers should be used with caution when dealing with sparsely connected groups. Empirical measures of the relative centrality across groups are extremely sensitive to measurement error in the links across groups. From the modeling perspective, the stark conclusions for investment and consumption decisions across groups also raise questions about the modeling assumptions. The usual justification for linear spillovers is that it is good local approximation to any continuos influence function. Our results point that this approximation causes some measures to be too sensitive to be useful for groups that are sparsely connected.

At the same time, our results highlight that centrality inside groups that are well connected can be robustly measured as it does not change with small perturbations across groups.

One of the most interesting directions for further work is to endogenize the network. Our approach views the network is exogenous, which is reasonable when constraints like languages, occupations, and geography determine much of the interaction that goes on in the short and medium run. In that case, the formation of the new links across groups arises exogenously because these parameters are varied by external forces. Nevertheless, over longer time scales agents do have choices in their interactions, though these are constrained and costly. The main challenge is to formulate a model which can encompass a fairly broad range of constraints on the network formation problem (such as different costs of linking between different pairs of agents) and produce tractable and interesting network formation dynamics.

References

Acemoglu, D., A. Ozdaglar, and A. ParandehGheibi (2009): "Spread of (Mis)Information in Social Networks," *Games and Economic Behavior*, forthcoming.

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

Bonachich, P. B. (1987): "Power and Centrality: A Family of Measures," *American Journal of Sociology*, 92, 1170–1182.

Bonacich, P. B. (1991): "Simultaneous Group and Individual Centralities," *Social Networks*, 13, 155–168.

Bramoullé, Y., R. Kranton, and M. D'Amours (2009): "Strategic Interaction and Networks," *Working Paper*.

Burt, R. S. (1992): *Structural Holes: The Social Structure of Competition*, Cambridge, MA: Harvard University Press.

Calvó-Armengol, A., E. Patacchini, and Y. Zenou (2009): "Peer Effects and Social Networks in Education," *Review of Economic Studies*, 76, 1239–1267.

Cameron, S. V. and J. J. Heckman (2001): "The Dynamics of Educational Attainment for Black, Hispanic, and White Males," *Journal of Political Economy*, 109, 455–499.

Conlisk, J. (1985): "Comparative Statics for Markov Chains," *Journal of Economic Dynamics and Control*, 9, 139–151.

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

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

Diestel, R. (2005): *Graph Theory*, vol. 173 of *Graduate Texts in Mathematics*, Springer-Verlag, Heidelberg, third ed.

Durrett, R. (2005): *Probability: Theory and Examples, Third Edition*, Belmont, CA: Thomson.

French, J. (1956): "A Formal Theory of Social Power," *Psychological Review*, 63, 181–194.

Golub, B. and M. Jackson (2009): "How Homophily affects Diffusion and Learning in Networks," *Working paper*.

Golub, B. and M. O. Jackson (2010): "Naïve learning in social networks: Convergence, influence, and the wisdom of crowds," *American Economic Journal: Microeconomics*, forthcoming.

Granovetter, M. (1973): "The Strength of Weak Ties," *The American Journal of Sociology*, 78, 1360–1380.

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

Katz, L. (1953): "A New Status Index Derived from Sociometric Analysis," *Psychometrika*, 18, 39–43.

Langville, A. N. and C. D. Meyer (2006): *Google's PageRank and Beyond: The Science of Search Engine Rankings.*, Princeton, N.J.: Princeton University Press.

Lever, C. (2010): "Voting competitions over networks: Whom to target?" *Working paper.*

McPherson, M., L. Smith-Lovin, and J. Cook (2001): "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology*, 27, 415–44.

Meyer, C. D. (2000): *Matrix Analysis and Applied Linear Algebra*, Philadelphia: SIAM.

Wasserman, S. and K. Faust (1994): *Social Network Analysis: Methods and Applications*, Cambridge: Cambridge University Press.

Appendix

Proof of Theorem 1

The proof amounts to checking that the claimed specification — call it $\widetilde{\mathbf{e}}(\mathbf{A})$ — satisfies the eigenvector condition

$$\widetilde{e}_i(\mathbf{A}) = \sum_j a_{ij}\widetilde{e}_j(\mathbf{A})$$

for every i. To do this, we use the facts that $\mathbf{e}(\mathbf{A}^X)$ is an eigenvector centrality (i.e. an eigenvector with eigenvalue 1) of $\mathbf{A}^X$, and $\mathbf{e}(\mathbf{A}^Y)$ is an eigenvector centrality of $\mathbf{A}^Y$. For example, to check the condition for $i = x$, we compute:

$$\begin{aligned}
\sum_j a_{xj}\widetilde{e}_j(\mathbf{A}) &= a_{xy}\widetilde{e}_y(\mathbf{A}) + a_{xx}\widetilde{e}_x(\mathbf{A}) + \sum_{i\in X\setminus\{x\}} a_{xi}\widetilde{e}_i(\mathbf{A}) \\
&= a_{xy}\cdot\frac{1}{1+r}\cdot\frac{1}{1-a_{xy}(1-e_y(\mathbf{A}^Y))}\cdot e_y(\mathbf{A}^Y) \\
&\quad + a_{xx}\cdot\frac{r}{1+r}\cdot\frac{1}{1-a_{yx}(1-e_x(\mathbf{A}^X))}\cdot e_x(\mathbf{A}^X) \\
&\quad + \sum_{i\in X\setminus\{x\}} a_{xi}\cdot\frac{r}{1+r}\cdot\frac{1-a_{yx}}{1-a_{yx}(1-e_x(\mathbf{A}^X))}\cdot e_i(\mathbf{A}^X)
\end{aligned}$$

$$
\begin{aligned}
&= a_{xy} \cdot \frac{1}{1+r} \cdot \frac{1}{1 - a_{xy}(1 - e_y(\mathbf{A}^Y))} \cdot e_y(\mathbf{A}^Y) \qquad \text{(Using the definition of } \mathbf{A}^X\text{)} \\
&\quad + a_{xx}^X (1 - a_{yx}) \cdot \frac{r}{1+r} \cdot \frac{1}{1 - a_{yx}(1 - e_x(\mathbf{A}^X))} \cdot e_x(\mathbf{A}^X) \\
&\quad + \sum_{i \in X \setminus \{x\}} a_{xi}^X \cdot \frac{r}{1+r} \cdot \frac{1 - a_{yx}}{1 - a_{yx}(1 - e_x(\mathbf{A}^X))} \cdot e_i(\mathbf{A}^X) \\
&= a_{xy} \cdot \frac{1}{1+r} \cdot \frac{1}{1 - a_{xy}(1 - e_y(\mathbf{A}^Y))} \cdot e_y(\mathbf{A}^Y) \\
&\quad + \frac{r}{1+r} \cdot \frac{1 - a_{yx}}{1 - a_{yx}(1 - e_x(\mathbf{A}^X))} \sum_{i \in X} a_{xi}^X e_i(\mathbf{A}^X) \\
&= a_{xy} \cdot \frac{1}{1+r} \cdot \frac{1}{1 - a_{xy}(1 - e_y(\mathbf{A}^Y))} \cdot e_y(\mathbf{A}^Y) \quad \text{(Using the definition of } \mathbf{e}(\mathbf{A}^X)\text{)} \\
&\quad + \frac{r}{1+r} \cdot \frac{1 - a_{yx}}{1 - a_{yx}(1 - e_x(\mathbf{A}^X))} \cdot e_x(\mathbf{A}^X)
\end{aligned}
$$

If we now plug in r from (1) and simplify, this is equal to

$$
\frac{r}{1+r} \cdot \frac{1}{1 - a_{yx}(1 - e_x(\mathbf{A}^X))} \cdot e_x(\mathbf{A}^X) = \widetilde{e}_x(\mathbf{A}).
$$

This verifies the eigenvector condition for $i = x$. The calculations for the other indices are equally straightforward.

Proof of Theorem 2

Part 1. If $\frac{\gamma}{c} < 1$ a Nash equilibrium exists and is unique.

Both of these statements follow from the same reasoning as Theorem 1 in Ballester, Calvó-Armengol, and Zenou (2006). The proof works by solving the first order conditions, which are a linear system of equations. The condition $\frac{\gamma}{c} < 1$ guarantees the linear system has a solution, given by $\mathbf{z}^*$ below, which proves existence.

$$
\mathbf{z}^* = \frac{1}{c}\mathbf{b}\left(\mathbf{A}, \frac{\gamma}{c}\right).
$$

To show uniqueness, note that payoffs are linear in the actions of the other players, so when opponents play mixed strategies, only the expectations of their choices matter. Additionally there is a unique best response to any mixed strategy because the cost of effort is convex. Therefore *all* equilibria in this regime are in pure strategies. When $\frac{\gamma}{c} < 1$, only $\mathbf{z}^*$ solves the first order conditions for a pure-strategy Nash equilibrium.

Part 2. If $\frac{\gamma}{c} \geq 1$, there is no pure-strategy Nash equilibrium.

Suppose, toward a contradiction, that $\alpha := \gamma/c \geq 1$ and there is an equilibrium $\mathbf{z}^*$. The first-order conditions imply that

$$\mathbf{z}^* = \alpha \mathbf{A}\mathbf{z}^* + \frac{1}{c}\mathbf{1},$$

By recursively substituting the entire right-hand side for $\mathbf{z}^*$ we get, for every natural number K:

$$\mathbf{z}^* = \frac{1}{c}\left[\sum_{k=0}^{K-1} \alpha^k \mathbf{A}^k\right]\mathbf{1} + \alpha^K \mathbf{A}^K \mathbf{z}.$$

Since $\mathbf{A}$ is column-stochastic, so is $\mathbf{A}^k$ for every k. In particular, column i of $\sum_{k=0}^{K} \alpha^k \mathbf{A}^k$ sums to $\sum_{k=0}^{K} \alpha^k \geq K$, and thus some entry of that column exceeds K/n. Choosing K so that $K/(cn)$ exceeds the maximum entry of $\mathbf{z}*$ yields a contradiction.

Part 3. If $\frac{\gamma}{c} \geq 1$, there is no Nash equilibrium.

Take a mixed-strategy Nash equilibrium $\mathbf{F}$, which is a vector of cumulative distribution functions. Create a pure strategy profile $\mathbf{z}$ where each j sets z_j to the expectation of his random investment under F_i. Using linearity of expectation as well as the fact that the payoff of i is linear in each z_j, we find that the function $z_i \mapsto u_i(z_i; \mathbf{z}_{-i})$ is the same as $z_i \mapsto u_i(z_i; \mathbf{F}_{-i})$, and so what was a best response remains a best response. This gives a pure-strategy Nash equilibrium. Thus, by the previous part, there cannot be a mixed-strategy Nash equilibrium in the regime $\frac{\gamma}{c} \geq 1$.

Proof of Theorem 3

Let $\|\cdot\|$ be the supremum norm on $\mathbb{R}^n$ and let $\|\cdot\|$ also denote the induced matrix norm when the argument is a matrix. Define $\mathbf{T} = \mu^{-1}\mathbf{M}$. Aperiodicity of $\mathbf{M}$ implies that $\mathbf{T}^k$ has a positive diagonal entry for some high enough k (Durrett, 2005, p. 310). That, in turn, implies that all eigenvalues of $\mathbf{A}$ are smaller than μ (Meyer, 2000, Section 8.3). From this it follows that $\mathbf{G} = \lim_{k\to\infty} \mathbf{T}^k$ is well-defined and that

$$\mathbf{G} = \frac{\mathbf{e}(\mathbf{T})\mathbf{e}(\mathbf{T}')'}{\mathbf{e}(\mathbf{T}')'\mathbf{e}(\mathbf{T})},$$

where $\mathbf{e}(\mathbf{T})$ is as in Definition 1 (Meyer, 2000, Section 8.3). The prime notation denotes transposition.

CLAIM 1. $\lim_{a\uparrow 1}(1-a)(\mathbf{I} - a\mathbf{T})^{-1}$ *exists and is equal to* $\mathbf{G}$.

Proof. Fix $\delta > 0$. Choose K so large that for $k > K$ we have $\left\|\mathbf{T}^k - \mathbf{G}\right\| < \delta/2$ and then choose $a < 1$ so that $\left|\sum_{k=0}^{K}(1-a)a^k\right| < \delta/4$. The Neumann series

$$(1-a)(\mathbf{I}-a\mathbf{T})^{-1} = \sum_{k=0}^{\infty}(1-a)a^k\mathbf{T}^k$$

converges (Meyer, 2000, Section 7.10). Now,

$$\begin{aligned}\left\|\sum_{k=0}^{K}(1-a)a^k\mathbf{T}^k - \mathbf{G}\right\| &= \left\|\sum_{k=0}^{\infty}(1-a)a^k\mathbf{T}^k - \sum_{k=0}^{\infty}(1-a)a^k\mathbf{G}\right\| \\ &\leq \sum_{k=0}^{K}(1-a)a^k\|\mathbf{T}^k - \mathbf{G}\| + \sum_{k=K+1}^{\infty}(1-a)a^k\|\mathbf{T}^k - \mathbf{G}\| \\ &\leq 2\sum_{k=0}^{K}(1-a)a^k + \sum_{k=K+1}^{\infty}(1-a)a^k\|\mathbf{T}^k - \mathbf{G}\| \\ &\leq 2\cdot\frac{\delta}{4} + \frac{\delta}{2}\sum_{k=K+1}^{\infty}(1-a)a^k \\ &\leq \delta.\end{aligned}$$

Here we have used the triangle inequality and the fact that $\mathbf{T}^k$ and $\mathbf{G}$ are both stochastic, so have matrix norm at most 1. □

Recall that that by Definition 3 we have

$$(1-a)\mathbf{b}(\mathbf{T},a) = (1-a)(\mathbf{I}-a\mathbf{T})^{-1}\mathbf{1}.$$

Thus, for any $\epsilon > 0$, there is a $\delta > 0$ so that the statement

$$\|(1-a)(\mathbf{I}-a\mathbf{T})^{-1} - \mathbf{G}\| < \delta$$

implies

$$\left|\frac{b_i(\mathbf{T},a)}{B(\mathbf{T},a)} - \frac{\sum_j G_{ij}}{\sum_{j,k} G_{jk}}\right| = \left|\frac{(1-a)b_i(\mathbf{T},a)}{(1-a)B(\mathbf{T},a)} - \frac{\sum_j G_{ij}}{\sum_{j,k} G_{jk}}\right| < \epsilon \text{ for all } i.$$

Recalling that

$$\frac{\sum_j G_{ij}}{\sum_{j,k} G_{jk}} = e_i(\mathbf{T})$$

and putting everything together with Claim 1 shows that for every $\epsilon > 0$, all high enough $a < 1$ satisfy

$$\left| \frac{b_i(\mathbf{T}, a)}{B(\mathbf{T}, a)} - \mathbf{e}(\mathbf{T}) \right| < \epsilon \text{ for all } i.$$

Since $\mathbf{b}(\mathbf{T}, a) = \mathbf{b}(\mathbf{M}, a\mu^{-1})$ and eigenvectors are invariant to scale, this completes the proof.

Proof of Theorem 4

We will show that the derivative of the logarithm of the ratio is positive.

First for a_{xy}:

$$\begin{aligned}
\frac{\partial \log\left(\frac{e_X(\mathbf{A})}{e_Y(\mathbf{A})}\right)}{\partial a_{xy}} &= \frac{\partial}{\partial a_{xy}}\Bigg(\log(a_{xy}) - \log(a_{yx}) + \log\left(\frac{e_y(\mathbf{A}^Y)}{e_x(\mathbf{A}^X)}\right) + \\
&\quad \log\left(1 - a_{yx}(1 - e_x(\mathbf{A}^X))\right) - \log\left(1 - a_{xy}(1 - e_y(\mathbf{A}^Y))\right)\Bigg) \\
&= \frac{1}{a_{xy}} + \frac{1 - e_y(\mathbf{A}^Y)}{1 - a_{xy}(1 - e_y(\mathbf{A}^Y))} > 0
\end{aligned}$$

Likewise for $e_y(\mathbf{A}^Y)$ we have

$$\begin{aligned}
\frac{\partial \log\left(\frac{e_X(\mathbf{A})}{e_Y(\mathbf{A})}\right)}{\partial e_y(\mathbf{A}^Y)} &= \frac{1}{e_y(\mathbf{A}^Y))} - \frac{a_{xy}}{1 - a_{xy}(1 - e_y(\mathbf{A}^Y))} \\
&= \frac{1 - a_{xy}}{e_y(\mathbf{A}^Y))\left(1 - a_{xy}(1 - e_y(\mathbf{A}^Y))\right)} > 0
\end{aligned}$$

The results for a_{yx} are $e_x(\mathbf{A}^X)$ follow by symmetry.

Proof of Theorem 5

Taking the first-order conditions for each agent's consumption decision we obtain:

$$\mathbf{z}^* = \frac{\beta}{1+\beta}\mathbf{A}\mathbf{z}^* + \frac{1}{1+\beta}\theta$$

Which is identical to the FOCs of the investment game with $\gamma = \beta$ and $c = 1 + \beta$. The rest of the proof follows from Theorem 2.

Proof of Corollary 4

This follows from Claim 1 of Theorem 3.