

A Network Approach to Public Goods

EXTENDED ABSTRACT

MATTHEW ELLIOTT AND BENJAMIN GOLUB[†]

Consider an economy in which agents ($N = \{1, 2, \dots, n\}$) each choose a level of effort, $a_i \geq 0$, toward an activity, such as mitigating pollution, that generates externalities. Agent i has a utility function $u_i : \mathbb{R}_+^n \rightarrow \mathbb{R}$ (where u_i is concave and continuously differentiable) mapping everyone's efforts to payoffs.

The incidence of externalities may be heterogeneous across those affected. How do the asymmetries and heterogeneities shape outcomes? The answers are very different depending on whether we consider noncooperative, voluntary provision or efficient negotiations. Prior work in networks has focused mostly on the first case: agents unilaterally decide how much to contribute. In that case, we can understand the scope for cooperation and the structure of contributions by looking at a network that represents the strategic spillovers in the environment: how one agent's actions affect others' marginal returns to contributing.¹

An important alternative benchmark to consider is an efficient solution, for instance one reached via multilateral negotiations. Our main thesis is that we can gain insight on such solutions by studying a network reflecting externalities. For any action profile \mathbf{a} under consideration, we construct a network in which the agents are nodes and the weighted links among them record the marginal externalities of actions. In particular, the link or arrow from agent i to j reflects how much marginal benefit i can confer on j by increasing i 's action. The structure of this network (in particular, its *spectral radius*) will help us diagnose the potential for Pareto improvements. In an important class of negotiated outcomes, agents' contributions are equal to their network centralities in a certain sense. The characterizations we obtain do not rely on parametric assumptions.

1 MAIN ASSUMPTIONS AND DEFINITIONS

The following two assumption capture that agents can, at a cost, provide nonrival positive externalities to each other.

Assumption 1 (Costly Actions (CA)). Each agent finds it costly to invest effort, holding others' actions fixed: $\frac{\partial u_i}{\partial a_i}(\mathbf{a}) < 0$ for any $\mathbf{a} \in \mathbb{R}_+^n$ and $i \in N$.

Assumption 2 (Positive Externalities (PE)). Increasing any agent's action level weakly benefits all other agents: $\frac{\partial u_i}{\partial a_j}(\mathbf{a}) \geq 0$ for any $\mathbf{a} \in \mathbb{R}_+^n$ whenever $j \neq i$.

Imagine for a moment that agents choose actions noncooperatively, playing a Nash equilibrium in the environment described. Costly Actions implies that the unique Nash equilibrium has $a_i = 0$ for each i . Except in the trivial case where the action profile $\mathbf{0}$ is already Pareto efficient, this means that there are Pareto gains to be realized by increasing actions together through some kind of structured favor-trading that improves on the static Nash outcome.²

[†]Elliott: Faculty of Economics, University of Cambridge, matthew.elliott@gmail.com. Golub: Department of Economics, Harvard, bgolub@fas.harvard.edu.

¹See, e.g., Ballester et al. [2006] and Bramoullé et al. [2014]

²One interpretation of the action profile $\mathbf{a} = \mathbf{0}$ is as a status quo at which negotiations begin. An alternative interpretation is that it is a Nash equilibrium in which everyone has already exhausted their private gains from exerting effort.

The Jacobian, $\mathbf{J}(\mathbf{a})$, is the n -by- n matrix whose (i, j) entry is $J_{ij}(\mathbf{a}) = \partial u_i(\mathbf{a}) / \partial a_j$. The *benefits matrix* $\mathbf{B}(\mathbf{a}; \mathbf{u})$ is then defined as follows:

$$B_{ij}(\mathbf{a}; \mathbf{u}) = \begin{cases} \frac{J_{ij}(\mathbf{a}; \mathbf{u})}{-J_{ii}(\mathbf{a}; \mathbf{u})} & \text{if } i \neq j \\ 0 & \text{otherwise.} \end{cases}$$

When $i \neq j$, the quantity $B_{ij}(\mathbf{a}; \mathbf{u})$ is i 's marginal rate of substitution between decreasing own effort and receiving help from j . In other words, it is how much i values the help of j , measured in the number of units of effort that i would be willing to put forth in order to receive one unit of j 's effort.

Since $J_{ii}(\mathbf{a}; \mathbf{u}) < 0$ by Assumption 1, the benefits matrix is well-defined. Assumptions 1 and 2 imply that it is entrywise nonnegative. We also impose:

Assumption 3 (Connectedness of Benefits (CONN)). For all $\mathbf{a} \in \mathbb{R}_+^n$, the matrix $\mathbf{B}(\mathbf{a})$ is irreducible.

This posits that it is not possible to find an outcome and partition society into two nonempty groups such that, at that outcome, one group does not derive any marginal benefit from the effort of the other group.

2 EFFICIENCY AND THE SPECTRAL RADIUS

This section shows that an important statistic of the benefits network—the size of the largest eigenvalue—can be used to diagnose whether an outcome is Pareto efficient. For any nonnegative matrix \mathbf{M} , we define $r(\mathbf{M})$ as the maximum of the magnitudes of the eigenvalues of \mathbf{M} , also called the *spectral radius*. This quantity can be interpreted as a single measure of how expansive a matrix is as a linear operator—how much it can scale up vectors that it acts on.

PROPOSITION 1.

- (i) Under Assumptions 1 (CA), 2 (PE), and 3 (CONN), an interior action profile $\mathbf{a} \in \mathbb{R}_{++}^n$ is Pareto efficient if and only if the spectral radius of $\mathbf{B}(\mathbf{a})$ is 1.
- (ii) Under Assumptions 1 (CA) and 2 (PE), the outcome $\mathbf{0}$ is Pareto efficient if and only if $r(\mathbf{B}(\mathbf{0})) \leq 1$.

We will discuss the proof of (i) only. The idea of the argument is that if the spectral radius is greater than 1, a Pareto improvement can be constructed in which one agent increases his action, generating benefits for others; then others “pass forward” some of the benefits they receive by increasing their own actions. Now we will actually construct such a Pareto improvement. We will fix any $\mathbf{a} \in \mathbb{R}_{++}^n$ (which we will not explicitly write as an argument everywhere) and let ρ denote the spectral radius of $\mathbf{B}(\mathbf{a})$. Then by the Perron-Frobenius Theorem and the maintained assumptions, there is a $\mathbf{d} \in \mathbb{R}_{++}^n$ such that $\mathbf{B}\mathbf{d} = \rho\mathbf{d}$. Multiplying each row of this matrix equation by $-J_{ii}(\mathbf{a})$, we find that for each i ,

$$\sum_{j \neq i} \frac{\partial u_i}{\partial a_j} d_j + \rho \frac{\partial u_i}{\partial a_i} d_i = 0.$$

If $\rho > 1$, then using the assumption of Costly Actions ($\frac{\partial u_i}{\partial a_i} < 0$) we deduce

$$\sum_{j \neq i} \frac{\partial u_i}{\partial a_j} d_j + \frac{\partial u_i}{\partial a_i} d_i > 0, \quad (1)$$

showing that a slight change where each i increases his action by the amount d_i yields a Pareto improvement. The vector \mathbf{d} describes the relative magnitudes of contributions to make the passing forward of benefits work out to achieve a Pareto improvement. Note that it is key to the argument that \mathbf{d} is positive. The conditions of the Perron-Frobenius Theorem guarantee the positivity of \mathbf{d} . If $\rho < 1$, we reason similarly to conclude the inequality (1) when we set \mathbf{d} to be *minus* the Perron vector of \mathbf{B} —i.e., when each i slightly *decreases* his action by the amount $|d_i|$.

Conversely, we now show that if $\rho(\mathbf{B}(\mathbf{a})) = 1$ then \mathbf{a} is Pareto efficient. To show this, note that by the Perron-Frobenius Theorem, the condition $\rho(\mathbf{B}(\mathbf{a})) = 1$ implies the existence of a *left-hand* eigenvector $\boldsymbol{\theta}$ of $\mathbf{B}(\mathbf{a})$, with all positive entries, satisfying $\boldsymbol{\theta}\mathbf{B}(\mathbf{a}) = \boldsymbol{\theta}$. This can readily be rearranged into the equation $\boldsymbol{\theta}\mathbf{J}(\mathbf{a}) = \mathbf{0}$, which is the system of first-order conditions for the problem of maximizing $\sum_i \theta_i u_i(\mathbf{a})$ by choosing \mathbf{a} . Since the first-order conditions hold for the vector of weights $\boldsymbol{\theta}$ and the maximization problem is concave, it follows that \mathbf{a} is Pareto efficient.

Note that θ_i is player i 's left-hand eigenvector centrality in the network $\mathbf{B}(\mathbf{a})$. Thus, *the Pareto weights supporting \mathbf{a} as a planner's optimal outcome are actually agent's eigenvector centralities in a suitably defined network.*

2.1 Application: Essential agents

Are there any agents that are essential to negotiations in our setting and, if so, how can we identify them? The efficiency result above suggests a simple approach to this question. Suppose for a moment that a given agent may be exogenously unable to take any action other than $a_i = 0$. How much does such an exclusion hurt the prospects for voluntary cooperation by the other agents?

Without agent i , the benefits matrix at the status quo of $\mathbf{0}$ is equal to the original $\mathbf{B}(\mathbf{0})$ with each entry in that row in i 's column set to 0. Call that matrix $\mathbf{B}^{[-i]}(\mathbf{0})$. Its spectral radius is no greater than that of $\mathbf{B}(\mathbf{0})$. In terms of consequences for efficiency, the most dramatic case is one in which the spectral radius of $\mathbf{B}(\mathbf{0})$ exceeds 1 but the spectral radius of $\mathbf{B}^{[-i]}(\mathbf{0})$ is less than 1. Then by Proposition 1(ii), a Pareto improvement on $\mathbf{0}$ exists when i is present but not when i is absent.

Thus *agent i 's participation is essential to achieving any Pareto improvement on the status quo precisely when his removal changes the spectral radius of the benefits matrix at the status quo from being greater than 1 to being less than 1.* To illustrate this, consider the following example in which $N = \{1, 2, 3, 4\}$.




Fig. 1. A graph corresponding to a benefits matrix in which agent #4 is essential despite providing smaller benefits than the others.

The import of the example is that agent 4, even though he confers the smallest marginal benefits, is the only essential agent.

Without him, there are no cycles at all and the spectral radius of the corresponding benefits matrix $\mathbf{B}^{[-4]}(\mathbf{0})$ is 0. On the other hand, when he is present but any one other agent ($i \neq 4$) is absent, then there is a cycle whose edges multiply to more than 1, and the spectral radius of $\mathbf{B}^{[-i]}(\mathbf{0})$ exceeds 1. Thus, the participation of a seemingly “small” agent in negotiations can make an essential difference to the ability to improve on the status quo when that agent completes cycles in the benefits network.

2.2 Spectral Radius in Terms of Cycles

A standard fact permits a general expression of the role cycles play in the spectral radius (for background and a proof, see, e.g., [Milnor 2001]).

Fact 1. For any nonnegative matrix \mathbf{M} , its spectral radius $r(\mathbf{M})$ is equal to $\limsup_{\ell \rightarrow \infty} \text{trace}(\mathbf{M}^\ell)^{1/\ell}$.

For a directed, unweighted adjacency matrix \mathbf{M} , the quantity $\text{trace}(\mathbf{M}^\ell)$ counts the number of cycles of length ℓ in the corresponding network. More generally, for an arbitrary matrix \mathbf{M} it measures the strength of all cycles of length ℓ by taking the product of the edge weights for each such cycle, and then summing these values over all such cycles. Thus, by Fact 1, the (suitably adjusted) total value of long cycles provides an asymptotically exact estimate of the spectral radius.

The essential agents discussed in the last section will be those that are present in sufficiently many of the high value cycles in the network. Relatedly, a single weak link in a cycle will dramatically reduce the value of that cycle. Thus networks with an imbalanced structure, in which it is rare for those agents who could confer large marginal benefits on others to be the beneficiaries of others' efforts, will have a lower spectral radius and there will be less scope for cooperation.

2.3 Application: Subdividing Negotiations

Consider an arbitrary Pareto efficient outcome \mathbf{a}^* that a planner would like to achieve. Suppose that the agents cannot negotiate in the full group (perhaps because a large negotiation is too costly) and are divided into two subsets, M and M^c ; \mathbf{a}^* is proposed to each. Then each group can contemplate deviations from \mathbf{a}^* that are Pareto-improving *for that group*. How cheaply can a planner incentivize agents to stay with the original outcome rather than deviate?

For simplicity, set $J_{ii}(\mathbf{a}) = -1$ for each i and all \mathbf{a} , and assume, for this subsection only, that the planner may use transfers of a numeraire that enters each agent's utility additively. New payoffs are $\tilde{u}_i(\mathbf{a}) = u_i(\mathbf{a}) + m_i(\mathbf{a})$, where $m_i(\mathbf{a})$ must be nonnegative. The profile $(m_i)_{i \in N}$ *deters deviations from \mathbf{a}^** if the restriction of \mathbf{a}^* to M is Pareto efficient for the population M with preferences $(\tilde{u}_i(\mathbf{a}))_{i \in N}$, and if the analogous statement holds for M^c . We care about bounding the *cost of separation* $c_M(\mathbf{a}^*)$, defined as the infimum of $\sum_{i \in N} m_i(\mathbf{a}^*)$ —payments made by the planner at the implemented outcome—taken over all profiles $(m_i)_{i \in N}$ that deter deviations from \mathbf{a}^* .

PROPOSITION 2.1. *Consider a Pareto efficient outcome \mathbf{a}^* , and let $\boldsymbol{\theta}$ be the corresponding Pareto weights. Then*

$$c_M(\mathbf{a}^*) \leq \sum_j \frac{\theta_i}{\theta_j} B_{ij}(\mathbf{a}^*) a_j^*,$$

where the summation is taken over all ordered pairs (i, j) such that one element is in M and the other is in M^c .

Holding \mathbf{a}^* fixed, the bound in the proposition is small when the network given by $\mathbf{B}(\mathbf{a}^*)$ has small total weight on links across groups—i.e. when the *cut* between them is small. Note that it is the properties of *marginal* benefits that are key. The proposition shows that a negotiation can be very efficiently separable even when the separated groups provide large total (i.e., inframarginal) benefits to each other. The question of when one can find a split with this property is discussed in a large literature in applied mathematics on spectral clustering and the spectral gap. One conclusion is that if there is an eigenvalue of $\mathbf{B}(\mathbf{a}^*)$ near its largest eigenvalue (1 in this case, since \mathbf{a}^* is efficient) then such a split exists [Hartfiel and Meyer 1998].

3 LINEAR FAVOR-TRADING: LINDAHL OUTCOMES

In this section, we focus attention on a particular class of Pareto efficient solutions. The idea is that a public good would be provided efficiently if we could replicate markets, allowing agents to trade favors at some exchange rate. Then they should be traded up to the point where the marginal social benefits of agents' efforts equal the marginal cost. Following this idea, we augment our setting with prices and look for a Walrasian equilibrium of the augmented economy.³ In the Online Appendix of the full paper, we discuss bargaining and implementation theory foundations for this solution in detail.

To define Lindahl outcomes, let \mathbf{P} be an n -by- n matrix of prices, with P_{ij} (for $i \neq j$) being the price i pays to j per unit of j 's effort. Let Q_{ij} be how much i purchases of j 's effort at this price. The total expenditure of i on other agents' efforts is $\sum_j P_{ij}Q_{ij}$ and the total income that i receives from other agents is $\sum_j P_{ji}Q_{ji}$. Market-clearing requires that all agents $i \neq j$ demand exactly the same effort from agent j , and so $Q_{ij} = a_j$ for all i and all $j \neq i$. Incorporating these market clearing conditions, agent i faces the budget constraint

$$\sum_{j:j \neq i} P_{ij}a_j \leq a_i \sum_{j:j \neq i} P_{ji}. \quad (\text{BB}_i(\mathbf{P}))$$

The Lindahl solution requires that, subject to market-clearing and budget constraints, the outcome is each agent's most preferred action profile among those he can afford:

Definition 1. An action profile \mathbf{a}^* is a *Lindahl outcome* for a preference profile \mathbf{u} if there are prices \mathbf{P} so that the following conditions hold for every i :

- (i) $\text{BB}_i(\mathbf{P})$ is satisfied when $\mathbf{a} = \mathbf{a}^*$;
- (ii) for any \mathbf{a} such that the inequality $\text{BB}_i(\mathbf{P})$ is satisfied, we have $u_i(\mathbf{a}^*) \geq u_i(\mathbf{a})$.

The main result in this section, Theorem 1, relates agents' contributions in Lindahl outcomes to how "central" they are in the network of externalities.

Definition 2. An action profile $\mathbf{a} \in \mathbb{R}_+^n$ has the *centrality property* (or is a *centrality action profile*) if $\mathbf{a} \neq \mathbf{0}$ and $\mathbf{B}(\mathbf{a})\mathbf{a} = \mathbf{a}$.

³A bargaining game studied by Dávila, Eeckhout, and Martinelli [2009] and Penta [2011] provides a noncooperative foundation for these outcomes. Agents take turns proposing, and each may say: "For every unit done by me, I demand that each agent j contribute r_j ." Following this, each agent simultaneously replies whether he vetoes the proposal, and if not, how many units he is willing to contribute at most. Assuming no vetoes, the maximum contributions are implemented consistent with the announced ratios and everyone's caps. If someone vetoes, the next proposer speaks.

According to this condition, \mathbf{a} is a right-hand eigenvector of $\mathbf{B}(\mathbf{a})$ with eigenvalue 1. The entries of this vector are widely studied as measures of the importance or centrality of nodes in a network. The centrality property says that, for each i , we have

$$a_i = \sum_{j \in N} B_{ij}a_j. \quad (2)$$

Equation (2) asserts that each agent's contribution is a weighted sum of the other agents' contributions, where the weight on a_j is proportional to the marginal benefits that j confers on i .

THEOREM 1. *The following are equivalent for a nonzero $\mathbf{a} \in \mathbb{R}_+^n$:*

- (i) $\mathbf{B}(\mathbf{a})\mathbf{a} = \mathbf{a}$, i.e., \mathbf{a} has the centrality property;
- (ii) \mathbf{a} is a Lindahl outcome.

Informally, Lindahl outcomes exist because they are Walrasian equilibria, and Walrasian equilibria exist for the usual reasons.⁴

3.1 Coalitional Deviations: A Core Property

As we are modeling negotiations, we can ask whether some subset of the agents could do better by breaking away and coming to some other agreement among themselves. The Lindahl outcomes turn out to deter such deviations, if we assume that following a deviation, the non-deviating agents revert to their unilateral best-response actions of 0.

Then we have the following result, which we formalize in the full paper: If $\mathbf{a} \in \mathbb{R}_+^n$ has the centrality property, then in any coalition M , some agent is worse off after any deviation from \mathbf{a} .

The remarkable yet simple argument for this, due to Shapley and Shubik [1969] is that the standard core of the artificial economy we presented earlier (with tradeable externalities) can be identified with the set of action profiles that are robust to coalitional deviations in our setting. Then the standard argument for the core property of Walrasian market outcomes applies. Details are carried out in the full paper.

3.2 An Outline of the Proof of Theorem 1

The key fact for the more difficult "if" part is that the system of equations $\mathbf{B}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{a}^*$ allows us to extract Pareto weights that support the outcome \mathbf{a}^* as efficient; using those Pareto weights and the Jacobian, we can construct prices that support \mathbf{a}^* as a Lindahl outcome.

Now in more detail: Suppose we have a nonzero \mathbf{a}^* so that $\mathbf{B}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{a}^*$. Assumption 3 (CONN) and the Perron-Frobenius theorem imply the profile \mathbf{a}^* is interior. The fact that $\mathbf{B}(\mathbf{a}^*)$ has 1 as an eigenvalue implies, as discussed in the proof of Proposition 1 above, the existence of a $\boldsymbol{\theta} \in \mathbb{R}_{++}^n$ such that $\boldsymbol{\theta}\mathbf{B}(\mathbf{a}^*) = \boldsymbol{\theta}$.

Let us normalize utility functions so that $J_{ii}(\mathbf{a}^*) = -1$. We will guess Lindahl prices

$$P_{ij} = \theta_i J_{ij}(\mathbf{a}^*) \text{ for } i \neq j.$$

For notational convenience, we also define a quantity $P_{ii} = \theta_i J_{ii}(\mathbf{a}^*)$. To show that at these prices, actions \mathbf{a}^* are a Lindahl outcome, two conditions must hold. The first is the budget-balance condition ($\text{BB}_i(\mathbf{P})$). Second, agents must be choosing optimal action levels subject to their budget constraints, given the prices.

⁴Formally, things are a little trickier: in our setting standard proofs do not go through because of their boundedness requirements, but in the longer paper we establish the existence of Lindahl outcomes generally, using our main characterization.

First, we will show that at the prices we've guessed, equation $BB_i(\mathbf{P})$ holds with equality and so each agent is exhausting his budget:

$$\sum_{j:j \neq i} P_{ij} a_j^* - a_i^* \sum_{j:j \neq i} P_{ji} = 0. \quad (3)$$

To this end, note that $\theta \mathbf{B}(\mathbf{a}^*) = \theta$ is equivalent to

$$\sum_{i \in N} \theta_i J_{ij}(\mathbf{a}^*) = 0 \quad \Leftrightarrow \quad \sum_{j: j \neq i} P_{ij} = -P_{ii},$$

where the rewriting on the right is from our definition of the P_{ij} . Now, the equation (3) that we would like to establish becomes

$$\sum_{j: j \neq i} P_{ij} a_j^* + a_i^* P_{ii} = 0$$

or $\mathbf{P} \mathbf{a}^* = \mathbf{0}$. Because row i of \mathbf{P} is a scaling of row i of $\mathbf{J}(\mathbf{a}^*)$, this is equivalent to $\mathbf{J}(\mathbf{a}^*) \mathbf{a}^* = \mathbf{0}$. But this is readily seen to be equivalent $\mathbf{B}(\mathbf{a}^*) \mathbf{a}^* = \mathbf{a}^*$, which we assumed.

It remains only to see that each agent is optimizing at prices \mathbf{P} . The essential reason for this is that price ratios are equal to marginal rates of substitution by construction. Indeed, when all the denominators involved are nonzero, we may write:

$$\frac{P_{ij}}{P_{ik}} = \frac{\theta_i J_{ij}(\mathbf{a}^*)}{\theta_i J_{ik}(\mathbf{a}^*)} = \frac{J_{ij}(\mathbf{a}^*)}{J_{ik}(\mathbf{a}^*)}. \quad (4)$$

Since P_{ii} is minus the income that agent i receives per unit of action, this checks that each agent is making an optimal effort-supply decision, in addition to trading off all other goods optimally.

The converse implication—that if \mathbf{a}^* is a nonzero Lindahl outcome, then $\mathbf{B}(\mathbf{a}^*) \mathbf{a}^* = \mathbf{a}^*$ —is much easier. A nonzero Lindahl outcome \mathbf{a}^* can be shown to be interior. Given this, and that agents are optimizing given prices, we have $P_{ij}/P_{ik} = J_{ij}(\mathbf{a}^*)/J_{ik}(\mathbf{a}^*)$, which echoes (4) above. In other words, each row of \mathbf{P} is a scaling of the same row of $\mathbf{J}(\mathbf{a}^*)$. Therefore, the condition that each agent is exhausting his budget,⁵ which can be succinctly written as $\mathbf{P} \mathbf{a}^* = \mathbf{0}$, implies that $\mathbf{J}(\mathbf{a}^*) \mathbf{a}^* = \mathbf{0}$, which is equivalent to $\mathbf{B}(\mathbf{a}^*) \mathbf{a}^* = \mathbf{a}^*$, as we've said above.

3.3 Explicit Formulas for Lindahl Outcomes

The eigenvalue and eigenvector conditions we have derived are implicit in that \mathbf{a} appears as an argument of \mathbf{B} in both Proposition 1 and Theorem 1. However, under specific functional forms, the latter characterization can be written as an explicit function of parameters of the model.

The preferences we consider are:

$$u_i(\mathbf{a}) = -a_i + \sum_j [\alpha G_{ij} a_j + H_{ij} \log a_j], \quad (5)$$

where \mathbf{G} and \mathbf{H} are nonnegative matrices (networks) with zeros on the diagonal (no self-links) and $\alpha < 1/r(\mathbf{G})$. Let $h_i = \sum_j H_{ij}$. For any preferences in this family, the centrality property ($\mathbf{a} = \mathbf{B}(\mathbf{a}) \mathbf{a}$) reduces to the equation $\mathbf{a} = \mathbf{h} + \alpha \mathbf{G} \mathbf{a}$.

If $\alpha = 0$, then $a_i = h_i$ and i 's Lindahl action is equal to the number of i 's neighbors in \mathbf{H} , which is i 's *degree centrality*. If, instead, $h_i = 1$ for all i , then $\mathbf{a} = [\mathbf{I} - \alpha \mathbf{G}]^{-1} \mathbf{1}$. This is the vector of *Bonacich centralities*, which are weighted sums of the numbers of walks of various lengths.⁶ Finally, as α approaches 1, agents'

⁵This follows because each agent is optimizing given prices, and by Assumption 3 there is always some contribution each agent wishes to purchase.

⁶For background, see Ballester, Calvó-Armengol, and Zenou [2006, Section 3] and [Jackson 2008, Section 2.2.4].

actions become proportional to their normalized eigenvector centralities in \mathbf{G} . As Lindahl outcomes are defined in terms of prices, the formulas we have presented may be viewed as microfoundations or interpretations of eigenvector centrality in price terms.

4 RELATED LITERATURE AND CONCLUSION

The interdependence or “networked” nature of economies is one of their defining features, and recent research has aimed to exploit network methods to understand them better. In addition to the literature on network games discussed in the introduction, a few articles have taken a network perspective in settings closer to ours. Perhaps the closest is Ghosh and Mahdian [2008], a model of negotiations in a linear environment. Agents benefit linearly from their neighbors' contributions, with a cap on how much each can contribute. Their main result is that there is an equilibrium of their game that achieves the maximum feasible contributions if and only if the largest eigenvalue of a network weight matrix is greater than one—a result generalized by our Proposition 1. Du, Lehrer, and Pauzner [2015] study a competitive exchange economy with particular parametric (Cobb-Douglas) preferences, and characterize its price equilibria in terms of a matrix describing the preferences, which in some ways resembles our Theorem 1.

The reliance of all these prior results on parametric assumptions leaves open the possibility that they are dependent on the functional forms. Our contribution is to show, without parametric assumptions, that outcomes of negotiations with externalities can be characterized exactly by eigenvalue or centrality properties. In doing this, we give a new economic angle on these important mathematical concepts. In the opposite direction, the connection offers a new way to approach economic questions with results on positive matrices. Sections 2.1 and 2.3 give some first ideas on how to use the connection, and raise many questions that we hope the techniques will help with.

REFERENCES

- BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): “Who's Who in Networks. Wanted: The Key Player,” *Econometrica*, 74, 1403–1417.
- BRAMOULLÉ, Y., R. KRANTON, AND M. D'AMOURS (2014): “Strategic interaction and networks,” *The American Economic Review*, 104, 898–930.
- DÁVILA, J., J. ECKHOUT, AND C. MARTINELLI (2009): “Bargaining Over Public Goods,” *Journal of Public Economic Theory*, 11, 927–945.
- DU, Y., E. LEHRER, AND A. PAUZNER (2015): “Competitive Economy as a Ranking Device over Networks,” *Games and Economic Behavior*, forthcoming.
- GHOSH, A. AND M. MAHDIAN (2008): “Charity Auctions on Social Networks,” in *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, 1019–1028.
- HARTFIEL, D. J. 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. O. (2008): *Social and Economic Networks*, Princeton, NJ: Princeton University Press.
- MILNOR, J. W. (2001): “Matrix Estimates and the Perron–Frobenius Theorem,” Notes, available at <http://www.math.sunysb.edu/~jack/DYNOTES/dnA.pdf>. Accessed November 3, 2012.
- PENTA, A. (2011): “Multilateral Bargaining and Walrasian Equilibrium,” *Journal of Mathematical Economics*, 47, 417–424.
- SHAPLEY, L. AND M. SHUBIK (1969): “On the Core of an Economic System with Externalities,” *The American Economic Review*, 59, 678–684.