

A Network Approach to Public Goods

Matthew Elliott

University of Cambridge

Benjamin Golub

Harvard University

Suppose that agents can exert costly effort that creates nonrival, heterogeneous benefits for each other. At each possible outcome, a weighted, directed network describing marginal externalities is defined. We show that Pareto efficient outcomes are those at which the largest eigenvalue of the network is 1. An important set of efficient solutions—Lindahl outcomes—are characterized by contributions being proportional to agents' eigenvector centralities in the network. The outcomes we focus on are motivated by negotiations. We apply the results to identify who is essential for Pareto improvements, how to efficiently subdivide negotiations, and whom to optimally add to a team.

When economic agents produce public goods, mitigate public bads, or more generally create externalities, the incidence of the externalities is often heterogeneous across those affected. A nation's economic policies—for example, implementing a fiscal stimulus, legislating environ-

We gratefully acknowledge the research support and hospitality of Microsoft Research New England. Elliott acknowledges financial support from National Science Foundation grant 1518941. Kevin He provided exceptional research assistance. The guidance of Nageeb Ali, Abhijit Banerjee, Matthew O. Jackson, Andy Skrzypacz, and Bob Wilson has been essential. For detailed comments on previous drafts, we thank Gabriel Carroll, Arun Chandrasekhar, Sylvain Chassang, Anil Jain, Juuso Toikka, Xiao Yu Wang, and Ariel Zucker. For helpful suggestions, we thank Daron Acemoglu, Yann Bramoullé, Federico Echenique, Glenn Ellison, Maryam Farboodi, Alex Frankel, Drew Fudenberg, Andrea Galeotti, Jerry Green, Hari Govindan, Sanjeev Goyal, Johannes Hörner, Scott Kominers, David Kreps, John Ledyard, Jacob

Electronically published March 25, 2019

[*Journal of Political Economy*, 2019, vol. 127, no. 2]

© 2019 by The University of Chicago. All rights reserved. 0022-3808/2019/12702-0006\$10.00

mental regulations, or reducing trade barriers—benefit foreign economies differently. Investments by a firm in research yield different spillovers for various producers and consumers. Cities' mitigation of pollution matters most for neighbors sharing the same environmental resources. And within a firm, an employee's efforts (e.g., toward team production) will benefit other employees to different degrees. How does heterogeneity in the incidence of externalities translate into outcomes? Which agents contribute the most and the least? Whose effort is particularly critical?

An active research program addresses these questions by modeling agents playing a Nash equilibrium of a one-shot public goods game, in which they unilaterally choose how much effort to put forth; see, for example, Ballester, Calvó-Armengol, and Zenou (2006), Bramoullé, Kranton, and D'Amours (2014), and Allouch (2015, 2017). These works model externalities via particular functional forms in which a network is a set of parameters. Links describe the pairs of players who directly affect each other's payoffs or incentives, as when two people collaborate on a project. The main results then characterize equilibrium effort levels via certain network statistics. Since these statistics are major subjects of study in their own right, the connection yields a rich set of intuitions as well as analytical techniques for comparative statics, identifying "key players," and various other policy analyses.¹

We argue that it is valuable to also study different classes of solutions in a public goods economy—ones motivated by negotiations—and, paralleling the results above, to understand how network properties matter for these solutions. The static Nash equilibrium is a benchmark relevant in cases where decisions are unilateral, with limited scope for repetition or commitment. Under this solution concept, agents do not internalize the externalities of their effort. Indeed, in a public goods game, players free-ride on the contributions of others, leading to a classic "tragedy of the commons" problem. The resulting inefficiencies can be substantial; in the context of problems like climate change, some argue that they are disastrous. In cases where large gains can be realized by improving on the unilateral benchmark, institutions arise precisely to foster multilateral cooperation. Global summits,² the World Trade Organization, re-

Leshno, Jon Levin, Mihai Manea, Eric Maskin, Vikram Manjunath, Muriel Niederle, Michael Ostrovsky, Antonio Penta, Michael Powell, Phil Reny, Larry Samuelson, Ilya Segal, Balazs Szentes, Harald Uhlig, Leeat Yariv, Alex Wolitzky, and many seminar participants. Finally, we thank four anonymous reviewers and the editor for extensive and constructive comments.

¹ There are many empirical applications of these results. See, e.g., Calvó-Armengol, Patacchini, and Zenou (2009) and Acemoglu, García-Jimeno, and Robinson (2015). Other theoretical papers that examine different issues related to the provision of public goods on networks include Bramoullé and Kranton (2007) and Galeotti and Goyal (2010).

² For example, it was at the Rio Earth Summit that the first international treaty on climate change was hammered out. There have been several other summits and associated climate change agreements since.

search consortia, and corporate team-building practices all aim to mitigate free riding by facilitating commitment. Therefore, rather than working with the static Nash equilibrium, this paper focuses on the complementary benchmark of Pareto efficient public goods provision in the presence of nonuniform externalities.

Our contribution is to show that taking a network perspective on the system of externalities sheds new light on efficient outcomes and the scope for efficient cooperation. First, we provide a new characterization of when Pareto improvements are possible, which relates such improvements to cycles of favor trading, quantified in a suitable way. Second, we characterize certain efficient solutions—the Lindahl outcomes, which have microfoundations in terms of negotiation games. Our results describe agents' contributions at these outcomes in terms of their positions in the network of externalities. The insights that the analysis generates can help address questions such as who should be given a seat at the negotiating table or admitted to a team. In contrast to the previous work mentioned above, our characterizations are nonparametric: a "network" representation of marginal externality incidence arises naturally from general utility functions. Finally, we provide new economic foundations and intuitions for statistics that are widely used to measure the centrality of agents in a network by relating these statistics to concepts such as Pareto weights and market prices.

I. Example and Roadmap

We now present the essentials of the model in a simplified example. Section II defines all the primitives formally in the general case. Each agent has a one-dimensional effort/action choice, $a_i \geq 0$; it is costly for an agent to provide effort, which yields positive, nonrival externalities for (some) others. For a concrete example, suppose that there are three towns—X, Y, and Z—located as shown in figure 1A, each generating air and water pollution during production. Because of the direction of the prevailing wind, the air pollution of a town affects only those east of it. A river flows westward, so Z's water pollution affects X but not Y, which is located away from the river.

Town i can forgo $a_i \geq 0$ units of production at a net cost of a dollar per unit, reducing its pollution and creating positive externalities for others affected by that pollution. The important part of this assumption is that the value of forgone production outweighs private environmental benefits; this assumes that the net private benefits of increasing effort have already been exhausted if they were present. Let $u_i(a_x, a_y, a_z)$ denote i 's payoff.

Suppose that the leaders of the towns attend a summit to try to agree on improvements that will benefit all of them. We focus on like-for-like agreements, in which agents trade favors by providing the public good of

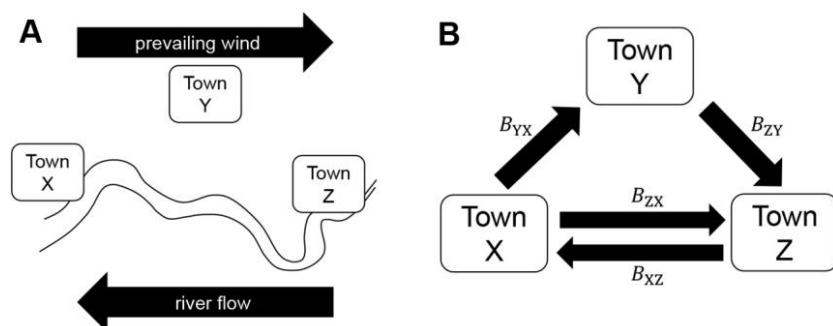


FIG. 1.—In this illustration of the framework, towns benefit from each other's pollution reduction. Town i benefits from j 's pollution reduction if pollution travels from j to i , which can happen via the wind or via the flow of the river. Let $B_{ij} = \partial u_i / \partial a_j$ be the marginal benefit to i from j 's reduction (per unit of i 's marginal cost, which is normalized to be 1). These numbers may vary with the action profile, (a_X, a_Y, a_Z) .

effort to each other, which is a relevant case for many practical negotiations.³ We begin by studying the set of all outcomes that are Pareto efficient and how they can be characterized in terms of the structure of externalities.

The conceptual platform for this—and for the rest of the paper—is to analyze a matrix whose entries record the marginal benefits per unit of marginal cost that each agent can confer on each other, for a given action profile. In our example, the entries of this matrix are

$$B_{ij}(\mathbf{a}) = \frac{\partial u_i}{\partial a_j}(\mathbf{a}) \bigg/ -\frac{\partial u_i}{\partial a_i}(\mathbf{a}) = \frac{\partial u_i}{\partial a_j}(\mathbf{a}),$$

for $i \neq j$, since we have normalized all marginal costs of effort to be 1. The diagonal terms of the matrix are set to 0, so that it records only the externalities between players and not their own costs. This *benefits matrix* can be equivalently represented as a (weighted, directed) network, where a link from i to j represents that i 's effort affects j 's welfare (see fig. 1B). That network is the key object whose statistics we relate to economic outcomes.

Our first result shows that an interior action profile \mathbf{a} is Pareto efficient if and only if 1 is a largest eigenvalue of $\mathbf{B}(\mathbf{a})$. The reason for this is as follows. The matrix $\mathbf{B}(\mathbf{a})$ is a linear system describing how investments translate into returns at the margin. Consider a particular sequence of investments. In figure 1B, Z can increase its action slightly and provide a marginal benefit to X. Then X, in turn, can “pass forward” some of the

³ This also parallels the above-mentioned papers regarding games on networks, which study one-dimensional contributions. In Sec. OA2 of the online appendix, we consider transfers: the very simple benchmark of quasi-linear preferences, as well as the general case, where our main results have natural analogues.

resulting increase in its utility, investing costly effort to help Z and Y. Finally, Y can also pass forward some of the increase in its utility by increasing its action, creating further benefits for Z. If they can all receive back more than they invest in such a multilateral adjustment, then the starting point is not Pareto efficient. It is in such cases that the linear system $\mathbf{B}(\mathbf{a})$ is “expansive”: there is scope for everyone to get more out of it than they put in. And an expansive system is characterized by having a largest eigenvalue exceeding 1. If the largest eigenvalue of $\mathbf{B}(\mathbf{a})$ is less than 1, then everyone can be made better off by reducing investment. As a result, the interior Pareto efficient outcomes have a benefits matrix with a largest eigenvalue exactly equal to 1. Section III.A makes this discussion rigorous (see proposition 1). Section III.B develops some of its interpretations and applications. It fleshes out the idea, already suggested by the informal argument, that cycles in the benefits network are critical for Pareto improvements and, correspondingly, that they determine the size of the largest eigenvalue. Finally, it discusses a simple algorithm to find the players who are essential to a negotiation—in the sense that without their participation, there is no Pareto improvement on the status quo. They are the ones whose removal causes a large disruption of cycles in the benefits network, as measured by the decrease in its largest eigenvalue.

One point on the Pareto frontier that is of particular interest is the classic Lindahl solution that completes the “missing markets” for externalities. If all externalities were instead tradable goods, we could consider the Walrasian outcome and identify the set of prices at which the market clears. If personalized taxes and subsidies equivalent to these prices could be charged in our public goods setting, then the same efficient outcome would obtain. Such an allocation is called a Lindahl outcome.⁴ Our second main result characterizes the Lindahl outcomes in terms of the *eigenvector centralities* of nodes in the marginal benefits network.

Eigenvector centrality is a way to impute importance to nodes in a network. Given a network \mathbf{G} , the eigenvector centrality of node i satisfies⁵

$$c_i \propto \sum_j G_{ij} c_j. \quad (1)$$

This equation says that i 's centrality is proportional to a weighted sum of its neighbors' centralities. Thus, the definition is a fixed-point condition and, in vector notation, becomes $\lambda \mathbf{c} = \mathbf{G}\mathbf{c}$ for some constant λ , so that the centralities of players are a right-hand eigenvector of the network \mathbf{G} . The measure captures the idea that central agents are those with strong connections to other central agents; equation (1) is simply a linear version of this statement. The notion of eigenvector centrality recurs in a

⁴ A formal definition of Lindahl outcomes appears in Sec. IV.

⁵ Under a network connectedness condition, these equations pin down relative centralities uniquely.

large variety of applications in various disciplines, and our main conceptual contribution is to relate it in a simple and general way to price equilibria. At the end of this section, we expand briefly on this point.

In our setting, we say that an action profile has the *centrality property* or, equivalently, is a *centrality action profile* if

$$\mathbf{a} = \mathbf{B}(\mathbf{a})\mathbf{a}. \quad (2)$$

We let the word “centrality” stand for eigenvector centrality and distinguish it explicitly from other kinds of centrality when necessary. Theorem 1 (in Sec. IV) establishes that Lindahl outcomes are exactly those with the centrality property. One way to interpret condition (2) is that i contributes in proportion to a weighted sum of others’ contributions a_j ; the weights are i ’s marginal valuations of the efforts of other agents.

Section V shows that the eigenvector condition (2) can be expressed in terms of walks in the benefits network, with the more central agents being those who sit at the locus of larger direct and indirect incoming marginal benefit flows. This relates price-based outcomes to the structure of the network. Building on this interpretation, theorem 1 is applied to study a problem in which a team has to decide which new member to admit. As another application, we study cases in which we can calculate the centrality action profiles explicitly. This, in turn, is used to give several important network centrality measures an economic microfoundation and interpretation in terms of price equilibria.⁶ This exercise echoes the general conceptual message of theorem 1—that there is a close connection between markets and network centrality—but for a wider range of network statistics and in a case where centralities can be computed explicitly in terms of exogenous parameters.

The Lindahl equilibria are of interest on more than just normative grounds. In Section IV.B, we review theories of negotiation that provide strategic foundations for this solution concept. First, using ideas from the literature on Walrasian bargaining (especially Dávila, Eeckhout, and Martinelli [2009] and Penta [2011]), we consider a model of multilateral negotiations that selects the Lindahl outcomes from the Pareto frontier. We then apply ideas of Hurwicz (1979a, 1979b) on implementation theory to show that the Lindahl equilibria are those selected by all mechanisms that are optimal in a certain sense. Finally, we note that, in our setting, Lindahl outcomes are robust to coalitional deviations—that is, are in an appropriately defined core.

We close by putting our work in a broader context of research on networks and centrality, beyond the most closely related papers on external-

⁶ Relatedly, Du, Lehrer, and Pauzner (2015) show how a ranking problem for locations on an unweighted graph can be studied via an associated perfectly competitive exchange economy in which agents have Cobb-Douglas utility functions. We discuss the connection in more detail in Sec. V.C.

ities and public goods. The interdependence of economic interactions is a defining feature of economies. When a firm does more business, it might employ more workers, who then have more income to spend on other goods, and so on. Eigenvector centrality (eq. [1]) loosely captures this idea. While in broad terms prior results suggest a connection between eigenvector centrality and economic outcomes, those results' reliance on parametric assumptions leaves open the possibility that the connection exists only in special cases and is heavily dependent on the functional forms. Our contribution is to show that the connection between centrality and markets goes deeper, by formalizing it in a simple model without parametric assumptions. In doing this, we give a new economic angle on a concept that has been the subject of much study. In sociology, key contributions on eigenvector-type centrality measures include Katz (1953), Bonacich (1987), and Friedkin (1991). For a survey of applications and results on network centrality from computer science and applied mathematics, especially for ranking problems, see Langville and Meyer (2012).⁷ Other applications include identifying those sectors in the macroeconomy that contribute the most to aggregate volatility via a network of intersectoral linkages (Acemoglu et al. 2012), analyzing communication in teams (Prat, de Martí, and Calvó-Armengol 2015), and the measurement of intellectual influence (Palacios-Huerta and Volij 2004). The last paper discusses axiomatic foundations of eigenvector centrality; other work taking an axiomatic approach includes Altman and Tennenholtz (2005) and Dequiedt and Zenou (2017).

We discuss other closely related literature in more detail at those points where we expect the comparisons to be most helpful. Omitted proofs and some supporting analyses are deferred to appendices.

II. Framework

A. *The Environment*

There is a set of agents or players, $N = \{1, 2, \dots, n\}$. The *outcome* is determined by specifying an action, $a_i \in \mathbb{R}_+$, for each agent i .⁸ Taking a higher action may be interpreted as doing more of something that helps the other agents—for instance, mitigating pollution. Agent i has a utility function $u_i : \mathbb{R}_+^n \rightarrow \mathbb{R}$, where u_i is concave and continuously differentiable; agent i 's payoff when the action profile \mathbf{a} is played is written $u_i(\mathbf{a})$.

⁷ Perhaps the most famous application of eigenvector centrality is the PageRank measure introduced as a part of Google's early algorithms to rank search results (Brin and Page 1998). For early antecedents of using eigenvectors as a way to "value" or rank nodes, see Wei (1952) and Kendall (1955).

⁸ We use \mathbb{R}_+ (\mathbb{R}_{++}) to denote the set of nonnegative (positive) real numbers. We write \mathbb{R}_+^n (\mathbb{R}_{++}^n) for the set of vectors \mathbf{v} with n entries such that each entry is in \mathbb{R}_+ (\mathbb{R}_{++}). When we write an inequality between vectors, e.g., $\mathbf{v} > \mathbf{w}$, that means that the inequality holds coordinate by coordinate; i.e., $v_i > w_i$ for each $i \in N$.

B. Main Assumptions

The following four assumptions are maintained in all results of the paper, unless a result explicitly states a different set of assumptions. Section VI.A discusses the extent to which some of our more economically restrictive explicit and implicit assumptions can be weakened.

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

Our results go through if efforts are required to be only weakly costly at the status quo. That allows us to interpret the status quo actions as an arbitrary Nash equilibrium of a game in which agents simultaneously choose how much effort to exert. We defer the technical issues associated with this generalization to Section V.B.

ASSUMPTION 2 (Positive externalities). Increasing any player's action level weakly benefits all other players: $\partial u_i(\mathbf{a})/\partial a_j \geq 0$ for any $\mathbf{a} \in \mathbb{R}_+^n$ whenever $j \neq i$.

Because the externalities are positive and nonrival, this is a public goods environment. Together, the two assumptions we have introduced—positive externalities and costly actions—imply that Pareto efficient outcomes will not be achieved if they are not equal to the status quo. The assumption of costly actions implies that the unique Nash equilibrium of a game in which players choose their actions entails that everyone contributes nothing ($a_i = 0$ for each i), even though other outcomes may Pareto dominate this one as a result of positive externalities—if those externalities are large enough.

One 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. We explore this second interpretation in Section V.B and, to accommodate it, relax assumption 1 to allow agents to have exactly zero net private benefits from increasing actions at the status quo.

Two additional technical assumptions are useful.

ASSUMPTION 3 (Connectedness of benefits). For all $\mathbf{a} \in \mathbb{R}_+^n$, if M is a nonempty proper subset of N , then there exist $i \in M$ and $j \notin M$ (which may depend on \mathbf{a}) such that $\partial u_i(\mathbf{a})/\partial a_j > 0$.

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.⁹

Finally, we assume that the set of points where everybody wants to scale up all effort levels is bounded. To state this, we introduce a few defini-

⁹ See Sec. OA6 of the online appendix for a discussion of extending the analysis when this assumption does not hold.

tions. Under a utility profile \mathbf{u} , action profile $\mathbf{a}' \in \mathbb{R}_+^n$ *Pareto dominates* another profile $\mathbf{a} \in \mathbb{R}_+^n$ if $u_i(\mathbf{a}') \geq u_i(\mathbf{a})$ for all $i \in N$, and the inequality is strict for some i . We say that \mathbf{a}' *strictly Pareto dominates* \mathbf{a} if $u_i(\mathbf{a}') > u_i(\mathbf{a})$ for all $i \in N$ and that \mathbf{a} is *Pareto efficient* (or simply *efficient*) if no other action profile Pareto dominates it.

ASSUMPTION 4 (Bounded improvements). The set

$$\{\mathbf{a} \in \mathbb{R}_+^n : \text{there is an } s > 1 \text{ so that } s\mathbf{a} \text{ strictly Pareto dominates } \mathbf{a}\}$$

is bounded.¹⁰

This assumption is necessary to keep the problem well behaved and ensure the existence of a Pareto frontier as well as of solutions to a bargaining problem we study below.¹¹ It is implied by, but weaker than, assuming an Inada condition whereby for high enough actions, marginal benefits become very low.

C. Key Notions

We write $\mathbf{u} = (u_1, u_2, \dots, u_n)$ for a profile of utility functions. The Jacobian, $\mathbf{J}(\mathbf{a}; \mathbf{u})$, is the n -by- n matrix whose (i, j) entry is $J_{ij}(\mathbf{a}; \mathbf{u}) = \partial u_i(\mathbf{a}) / \partial a_j$. The *benefits matrix* $\mathbf{B}(\mathbf{a}; \mathbf{u})$ is then defined as follows:

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

As discussed in the roadmap, 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.

Suppose that \mathbf{u} satisfies the assumptions of Section II.B. 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. Assumption 3 is equivalent to the statement that this matrix is irreducible at every \mathbf{a} .¹²

In discussing both the Jacobian and the benefits matrix, when there is no ambiguity about what \mathbf{u} is, we suppress it.

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*. That is,

¹⁰ This condition is weaker than assuming that the set of Pareto efficient outcomes is bounded.

¹¹ For details, see Sec. IV.B and particularly the proof of proposition 2.

¹² A matrix \mathbf{M} is *irreducible* if it is not possible to find a nonempty proper subset S of indices so that $M_{ij} = 0$ for every $i \in S$ and $j \notin S$.

$$r(\mathbf{M}) = \max\{|\lambda| : \lambda \text{ is an eigenvalue of } \mathbf{M}\},$$

where $|\lambda|$ denotes the absolute value of the complex number λ . By the Perron-Frobenius theorem (see app. A for a statement), any such matrix has a real, positive eigenvalue equal to $r(\mathbf{M})$. Thus, we may equivalently think of $r(\mathbf{M})$ as the largest eigenvalue of \mathbf{M} on the real line.

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. When applied to the benefits matrix \mathbf{B} , it will identify the scope for Pareto improvements.

III. Efficiency and the Spectral Radius

The thesis of this paper is that we can gain insight about efficient solutions to public goods problems by constructing, for any action profile \mathbf{a} under consideration, a network in which the agents are nodes and the weighted links among them measure the marginal benefits available by increasing actions. The adjacency matrix of this network is $\mathbf{B}(\mathbf{a})$.

This section offers support for the thesis by showing that an important statistic of this network—the size of the largest eigenvalue—can be used to diagnose whether an outcome is Pareto efficient (Sec. III.A). After presenting this general result, we discuss interpretations (especially in terms of the structure of the network) and applications.

A. A Characterization of Pareto Efficiency

Our main result on efficiency is the following.

PROPOSITION 1. (a) Under assumptions 1–3, 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. (b) Under assumptions 1 and 2, the outcome $\mathbf{0}$ is Pareto efficient if and only if $r(\mathbf{B}(\mathbf{0})) \leq 1$.

One argument for this result makes precise the intuition presented in the roadmap: when the spectral radius is greater than 1, we can obtain a Pareto improvement if one agent increases his action, generating benefits for others, and then other agents “pass forward” some of the benefits they receive. For simplicity, we work with claim a in the proposition, deferring claim b to the proof in appendix C. Fix any $\mathbf{a} \in \mathbb{R}_{++}^n$ and drop it in arguments; write ρ for 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 ($\partial u_i / \partial a_i < 0$), we deduce that

$$\sum_{j \neq i} \frac{\partial u_i}{\partial a_j} d_j + \frac{\partial u_i}{\partial a_i} d_i > 0, \quad (3)$$

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.¹³ The conditions of the Perron-Frobenius theorem guarantee the positivity of \mathbf{d} , though weaker conditions are known—see Section VI.A. If $\rho < 1$, we reason similarly to conclude the inequality (3) when each i slightly decreases his action by the amount d_i .

The key step not shown by the argument so far is 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 θ of $\mathbf{B}(\mathbf{a})$ with all positive entries, satisfying $\theta \mathbf{B}(\mathbf{a}) = \theta$. This can readily be rearranged into the equation $\theta \mathbf{J}(\mathbf{a}) = \mathbf{0}$, 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 θ and the maximization problem is concave, it follows that \mathbf{a} is Pareto efficient.

A complete proof of proposition 1 is in appendix C.

Proposition 1 shows that we can diagnose whether an outcome is Pareto efficient by using just the spectral radius of the benefits matrix, dispensing with the construct of Pareto weights. Moreover, the spectral radius provides more than just qualitative information; it can also be interpreted as a quantitative measure of the size of the inefficiency. In particular, the spectral radius measures the best return on investment in public goods per unit of cost that can simultaneously be achieved for all agents. Details on this can be found in appendix B. We axiomatize the spectral radius of the benefits matrix as a measure of marginal (in)efficiency in a sister paper, Elliott and Golub (2015).

The condition $\theta \mathbf{B}(\mathbf{a}^*) = \theta$ says that, for each i , we have $\theta_i = \sum_j \theta_j B_{ji}$. That is, i 's Pareto weight is equal to the sum of the various other Pareto weights, with θ_j weighted by $B_{ji}(\mathbf{a}^*)$, which measures how much j cares about i 's contribution. This echoes the definition of eigenvector centrality from Section I; indeed, θ is the eigenvector centrality of the network $\mathbf{B}(\mathbf{a}^*)^\top$. Thus, a planner maximizes the weighted sum of utilities, with weights θ , by having the agents take actions so that in the transpose of the induced benefits network each agent's centrality is equal to his Pareto

¹³ If its entries had different signs, then $\mathbf{B}\mathbf{d} - \rho\mathbf{d} > 0$ would not imply anything useful about Pareto improvements, because the second term would not move uniformly in one direction when ρ is replaced by 1.

weight. Correspondingly, in the transpose of the benefits network at a Pareto efficient outcome, each agent's centrality reveals his implied weighting by a planner.

The condition that the spectral radius of $\mathbf{B}(\mathbf{a})$ is 1 is independent of how different players' cardinal utilities are measured—as, of course, it must be, since Pareto efficiency is an ordinal notion. To see how the benefits matrix changes under reparameterizations of cardinal utility, suppose that we define, for each $i \in N$, new utility functions $\hat{u}_i(\mathbf{a}) = f_i(u_i(\mathbf{a}))$ for some differentiable, strictly increasing functions f_i . If we let $\hat{\mathbf{B}}$ be the benefits matrix obtained from these new utility functions, then $\mathbf{B}(\mathbf{a}) = \hat{\mathbf{B}}(\mathbf{a})$; this follows by applying the chain rule to the numerator and denominator in the definition of the benefits matrix.

Relation to utilitarian surplus.—It is worth comparing our condition for Pareto efficiency to the condition for efficiency from the perspective of a utilitarian planner who places equal weights on all players. Assume for this discussion that $\partial u_i / \partial a_i = -1$, so that all agents' utilities are normalized to make the costs of contribution equal in utils. Then, there is a utilitarian improvement if and only if some column of $\mathbf{B}(\mathbf{a})$ has a sum not equal to 1.¹⁴ This condition is very different from the spectral radius being different from 1, though utilitarian efficiency (all column sums equaling 1) does, of course, imply Pareto efficiency (the spectral radius equaling 1).

Some further general statements can be made, however. If all the column sums of $\mathbf{B}(\mathbf{a})$ are less than 1 (the utilitarian marginal surplus of every agent's contribution is less than 1), then the spectral radius of $\mathbf{B}(\mathbf{a})$ is less than 1 (Takayama 1985, corollary to theorem 4.C.11) and a Pareto improvement can be achieved. Similarly, if all column sums of $\mathbf{B}(\mathbf{a})$ exceed 1, then the spectral radius of $\mathbf{B}(\mathbf{a})$ exceeds 1 and there is also guaranteed to be a Pareto improvement. However, when some columns of $\mathbf{B}(\mathbf{a})$ have sums exceeding 1 and others do not, there is not a simple relation between these sums and Pareto improvements or even the possibility of Pareto improvements. We discuss these issues further in Elliott and Golub (2015).

B. Essential Players

Are there any players who are essential to negotiations in our setting, and if so, how can we identify them?

The efficiency results of Section III.A suggest a simple way of characterizing how essential any given player is to the negotiations. Suppose for

¹⁴ If the sum of column j exceeds 1, the utilitarian marginal social surplus of j increasing his contribution, $\sum_{i \neq j} \partial u_i / \partial a_j$, exceeds the social cost, which is 1. The argument for the other inequality is similar.

a moment that a given player exogenously may or may not be able to participate in an institution to negotiate an outcome that Pareto dominates the status quo. If he is not able, then his action is set to the status quo level of $a_i = 0$. How much does such an exclusion hurt the prospects for cooperation by the other agents?

Without player i , the benefits matrix at the status quo of $\mathbf{0}$ is equal to the original $\mathbf{B}(\mathbf{0})$ without row and column i ; equivalently, each entry in that row and column may be set to zero. Call a matrix constructed that way $\mathbf{B}^{[-i]}(\mathbf{0})$. The spectral radius of $\mathbf{B}^{[-i]}(\mathbf{0})$ 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 1b, a Pareto improvement on $\mathbf{0}$ exists when i is present but not when i is absent.

This argument shows that all players weakly improve the scope for Pareto improvements and that a player 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 directly apply this result involves calculating the spectral radius of many counterfactual benefits matrices. Is there a way in which we instead identify essential players by simply eyeballing the benefits network?

As noted in the roadmap, a cycle of players such that each can help the next creates scope for cooperation. When there are no cycles of cooperation at the status quo actions, there is no way to simultaneously compensate all members of any set of agents for taking positive effort, and no Pareto improvements are possible. Such a situation corresponds to the benefits matrix having a spectral radius of 0 at the status quo actions, and so the lack of cycles is directly tied to the spectral radius. Thus, a sufficient condition for a player to be essential is for that player to be part of all cycles.

To illustrate this, consider the following example illustrated in figure 2, in which $N = \{1, 2, 3, 4\}$.

The import of the example is that player 4, even though he confers the smallest marginal benefits, is the only essential player. 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 player ($i \neq 4$) 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" player in negotiations can make an essential difference to the ability to improve on the status quo when that player completes cycles in the benefits network.

The example suggests that we might be able to reinterpret proposition 1 in terms of the cycles that are present in the benefits matrix. However, the example is particularly stark—player 4 is involved in all cycles.




FIG. 2.—Example of a benefits matrix $\mathbf{B}(\mathbf{0})$ and its graphical depiction, in which player 4 is essential despite providing smaller benefits than the others.

More generally, how do different cycles feed into the spectral radius, and can we use that connection to identify essential players? A standard fact permits a general and useful interpretation (for background and a proof, see, e.g., Milnor [2001]).

FACT 1. (a) For any nonnegative matrix \mathbf{M} , $r(\mathbf{M}) = \limsup_{\ell \rightarrow \infty} \text{trace}(\mathbf{M}^\ell)$. (b) In particular, if \mathbf{B} and $\hat{\mathbf{B}}$ are two nonnegative matrices such that $\mathbf{B} \geq \hat{\mathbf{B}}$, then $r(\mathbf{B}) \geq r(\hat{\mathbf{B}})$.

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.¹⁵ Thus, by fact 1, the total value of long cycles provides an asymptotically exact estimate of the spectral radius.

An immediate implication of fact 1 is that essential players are those who are present in sufficiently many of the high-value cycles in the network, regardless of the specific marginal benefits they receive and provide. Ballester et al. (2006) pose a question similar to ours. They consider a setting where agents can privately benefit from taking positive effort and players simultaneously choose how much effort to exert. Studying the Nash equilibrium of this game, they define the key players as those whose removal results in the largest decrease in aggregate effort. In appendix E, we provide an example in which their key player and our essential player differ.¹⁶ Loosely, the essential player is the player present in many strong cycles of marginal benefits (measured by first derivatives of payoffs), while

¹⁵ More formally, a (directed) cycle of length ℓ in the matrix \mathbf{M} is a sequence $(c(1), c(2), \dots, c(\ell), c(\ell + 1))$ of elements of N (players), so that the cycle starts and ends at the same node ($c(\ell + 1) = c(1)$), and $M_{c(t)c(t+1)} > 0$ for each $t \in \{1, \dots, \ell\}$. Let $\mathcal{C}(\ell; \mathbf{M})$ be the set of all cycles of length ℓ in matrix \mathbf{M} . For any nonnegative matrix \mathbf{M} , $\text{trace}(\mathbf{M}^\ell) = \sum_{c \in \mathcal{C}(\ell; \mathbf{M})} \prod_{t=1}^{\ell} M_{c(t)c(t+1)}$.

¹⁶ Sec. V.B discusses the details of endogenizing the Nash status quo, which permits studying its comparative statics simultaneously with those of our efficient solution in the same model.

the key player is the player whose effort is most directly and indirectly complementary to others' (measured by cross partials of payoffs).

The connection between the spectral radius and cycles also suggests when there will be greater scope for cooperation. 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.

IV. Lindahl Outcomes

In this section, we focus attention on a particular class of Pareto efficient solutions. The insight behind the Lindahl solution is that a public good would be provided efficiently if each agent could be made to face a personalized price equal to his marginal benefit from the good. This would allow contributions to be collected up to the point where the marginal social benefit of providing the public good equals its marginal social cost. This point was initially made in simple environments, but Arrow (1969) shows that, quite generally, externalities—whatever their incidence—can be reinterpreted as missing markets. Following Lindahl and Arrow, we augment our setting by adding the missing markets and look for a Walrasian equilibrium of the augmented economy. We refer to these outcomes as Lindahl outcomes. The prices in the markets that are introduced are personalized taxes and subsidies: each agent pays a personalized tax for every unit of each other agent's effort he enjoys, and he receives a personalized subsidy (financed by others' taxes) per unit of effort he exerts.¹⁷ These prices are not subject to the normal equilibrating forces that operate in competitive markets (Samuelson 1954). In Section IV.B, we review game-theoretic microfoundations for the Lindahl concept in our setting, explaining what sorts of negotiations can lead to Lindahl outcomes.

To construct the augmented economy, 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

¹⁷ There need not be any transferable private commodity in which these prices are denominated. We can think of each player having access to artificial tokens, facing prices for the public goods denominated in these tokens, and being able to choose any outcome subject to not using more tokens than he receives from others.

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

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. We therefore have the following definition.

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

Hatfield et al. (2013) consider the problem of agents located on a network trading bilateral contracts. The augmented economy we have constructed can be mapped into their very general domain. They show that with quasi-linear utilities and under a condition of “full substitutability,” stable outcomes exist and are essentially equivalent to the competitive equilibrium outcomes. It might be hoped that we can make use of their results. Unfortunately, we cannot. Their full-substitutability condition is violated by our augmented economy. Intuitively, the opportunities for agent i to be compensated for his effort by agents j and k are complementary—agent i needs to exert effort only once to be compensated by both j and k .

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 $\mathbf{B}(\mathbf{a})\mathbf{a} = \mathbf{a}$.

According to this condition, \mathbf{a} is a right-hand eigenvector of $\mathbf{B}(\mathbf{a})$ with eigenvalue 1. Because actions are nonnegative, the Perron-Frobenius theorem implies that such an \mathbf{a} is the Perron, or principal, eigenvector—the one associated with the largest eigenvalue of the matrix.¹⁸ Section I provided some background on this notion of centrality.

THEOREM 1. The following are equivalent for a nonzero $\mathbf{a} \in \mathbb{R}_+^n$: (a) $\mathbf{B}(\mathbf{a})\mathbf{a} = \mathbf{a}$, that is, \mathbf{a} has the centrality property; and (b) \mathbf{a} is a Lindahl outcome.

We can also establish that any nonzero Lindahl outcome is interior (lemma 2; app. C). An outline of the proof of theorem 1 is below, and the complete proof appears in appendix C. However, before we present the argument, it is worth remarking on some simple consequences of

¹⁸ We discussed in Sec. III.A that the set of Pareto efficient action profiles is invariant to rescaling the utility functions, because such rescalings do not affect the benefits matrix. The same argument implies that the centrality action profiles are also invariant to such rescalings.

theorem 1. First, at any interior Lindahl outcome \mathbf{a} , the matrix $\mathbf{B}(\mathbf{a})$ has a nonnegative right eigenvector \mathbf{a} with eigenvalue 1 and therefore, by the Perron-Frobenius theorem, a spectral radius of 1. Proposition 1 then implies the Pareto efficiency of \mathbf{a} , providing an alternative proof of the first welfare theorem.¹⁹

Second, the condition $\mathbf{B}(\mathbf{a})\mathbf{a} = \mathbf{a}$ is a system of n equations in n unknowns (the coordinates of \mathbf{a}). By a standard argument (see, e.g., Shannon 2008), this entails that for generic utility functions satisfying our assumptions, the set of solutions will be of dimension 0 in \mathbb{R}_+^n . Therefore, the set of Lindahl outcomes is typically “small,” as is usually the case with sets of market equilibria.

Finally, the equivalence between Lindahl outcomes and centrality action profiles allows us to establish the existence of a Lindahl equilibrium in our setting, where standard proofs do not go through because of their boundedness requirements.

PROPOSITION 2. Either $\mathbf{a} = \mathbf{0}$ is Pareto efficient or there is a centrality action profile in which all actions are strictly positive.

The proof of proposition 2 is in appendix C. We also show that the profile $\mathbf{0}$ is a Lindahl outcome if and only if it is Pareto efficient (proposition 7; app. D).

A. An Outline of the Proof of Theorem 1

It is convenient to introduce scaling-indifferent action profiles. From the definition of the benefits matrix, scaling indifference is easily verified to be equivalent to the centrality property, and we use the two notions interchangeably. Recall that $\mathbf{J}(\mathbf{a})$ is the Jacobian, with entry (i, j) equal to $J_{ij}(\mathbf{a}) = \partial u_i(\mathbf{a}) / \partial a_j$.

DEFINITION 3. An action profile $\mathbf{a} \in \mathbb{R}_+^n$ satisfies *scaling indifference* (or *is scaling indifferent*)²⁰ if $\mathbf{a} \neq \mathbf{0}$ and $\mathbf{J}(\mathbf{a})\mathbf{a} = \mathbf{0}$.

We show below that a profile is a Lindahl outcome if and only if it has the centrality property. The more difficult part is the “if” part. The key fact 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, and using those Pareto weights and the Jacobian, we can construct prices that support \mathbf{a}^* as a Lindahl outcome.

Now we add more detail. Suppose that we have a nonzero \mathbf{a}^* so that $\mathbf{B}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{a}^*$. As we noted in the previous subsection, the profile \mathbf{a}^* is

¹⁹ A standard proof can be found in, e.g., Foley (1970).

²⁰ To see the reason for the name, note that, to a first-order approximation, $\mathbf{u}(\mathbf{a} + \varepsilon \mathbf{v}) \approx \mathbf{u}(\mathbf{a}) + \varepsilon \mathbf{J}(\mathbf{a})\mathbf{v}$. Suppose now that actions \mathbf{a} are scaled by $1 + \varepsilon$, for some small real number ε ; this corresponds to setting $\mathbf{v} = \mathbf{a}$. If $\mathbf{J}(\mathbf{a})\mathbf{a} = \mathbf{0}$, then all players are indifferent, in the first-order sense, to this small proportional perturbation in everyone’s actions.

then interior and Pareto efficient.²¹ It follows by a standard fact that there are Pareto weights $\theta \in \mathbb{R}_+ \setminus \{0\}$ such that \mathbf{a}^* maximizes $\sum_i \theta_i u_i(\mathbf{a})$ over all $\mathbf{a} \in \mathbb{R}_+^n$.

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

$$P_{ij} = \theta_i J_{ij}(\mathbf{a}^*)$$

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}^* satisfy the definition of a Lindahl outcome, two conditions must hold. The first is the budget balance condition, replicated for convenience:

$$\sum_{j:j \neq i} P_{ij} a_j^* - a_i^* \sum_{j:j \neq i} P_{ji} \leq 0. \quad (\text{BB}_i(\mathbf{P}'))$$

Second, agents must be choosing optimal action levels subject to their budget constraints, given the prices.

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

$$\sum_{j:j \neq i} P_{ij} a_j^* - a_i^* \sum_{j:j \neq i} P_{ji} = 0. \quad (4)$$

To this end, first note that \mathbf{a}^* maximizes $\sum_i \theta_i u_i(\mathbf{a})$, implying the first-order conditions

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

where the rewriting on the right is from our definition of the P_{ij} . Now, the equation (4) 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 $\mathbf{J}(\mathbf{a}^*) \mathbf{a}^* = \mathbf{0}$. This holds because \mathbf{a}^* is a centrality action profile and thus, as we noted at the beginning of this proof outline, scaling indifferent.

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

$$\frac{P_{ij}}{P_{ik}} = \frac{\theta_i J_{ij}(\mathbf{a}^*)}{\theta_i J_{ik}(\mathbf{a}^*)} = \frac{J_{ij}(\mathbf{a}^*)}{J_{ik}(\mathbf{a}^*)}. \quad (5)$$

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.

²¹ This is the point where the Perron-Frobenius theorem plays a key role—recall the discussion that follows proposition 1a.

Consider now the converse implication—that if \mathbf{a}^* is a nonzero Lindahl outcome, then $\mathbf{J}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{0}$. A nonzero Lindahl outcome \mathbf{a}^* can be shown to be interior. (This is lemma 2 in app. C.) Given this and that agents are optimizing given prices, we have

$$\frac{P_{ij}}{P_{ik}} = \frac{J_{ij}(\mathbf{a}^*)}{J_{ik}(\mathbf{a}^*)},$$

which echoes equation (5) 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,²² which can be succinctly written as $\mathbf{P}\mathbf{a}^* = \mathbf{0}$, implies that $\mathbf{J}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{0}$.

For intuition, we offer a brief comment on the form of the prices. The prices we guessed were $P_{ij} = \theta_i J_{ij}(\mathbf{a}^*)$. This entails that, all else equal, an agent pays a higher price if his Pareto weight is greater and if he values the good in question more (relative to his own marginal cost of providing effort—remembering that we have normalized here so that $J_{ii}(\mathbf{a}^*) = -1$). What is the reason for prices to have this form?

For agent i to be optimizing, he must be maximizing $u_i(\mathbf{a})$ subject to the budget constraint $(\text{BB}_i(\mathbf{P})')$ and, by the first-order conditions, $\mu_i P_{ij} = J_{ij}(\mathbf{a}^*)$, where μ_i is the Lagrange multiplier on the constraint $(\text{BB}_i(\mathbf{P})')$ —that is, the marginal utility of relaxing the constraint $(\text{BB}_i(\mathbf{P})')$, or the marginal utility of income to i . Next, consider the planner who puts weight θ_i on player i . At a solution to this planner's problem, it must be that $\mu_i \theta_i$ is the same across agents and thus a constant—otherwise the planner would want to increase the actions of some agents and reduce the actions of others. Combining these two observations, we deduce that P_{ij} is directly proportional to $\theta_i J_{ij}(\mathbf{a}^*)$, and, as only relative prices matter, we can set $P_{ij} = \theta_i J_{ij}(\mathbf{a}^*)$, which is the guess we made above.²³

B. A Review of Foundations for the Lindahl Solution

We introduced the Lindahl solution as a conceptual device for emulating missing markets for externalities but deferred discussion of how it can be implemented in actual negotiations over public goods. In this section, we review several foundations for the Lindahl solution—combinations of normative and strategic properties implying it. In view of theorem 1, these results are equivalently foundations for the class of centrality action profiles. Our discussions here adapt existing results, and so we describe the essence of each foundation briefly, referring to the prior literature. In each case, we have to adjust previous arguments to work in our setting

²² This follows because each agent is optimizing given prices, and by assumption 3 there is always some contribution each agent wishes to purchase.

²³ We thank Phil Reny for this insight.

with unbounded action spaces. Sections OA3–OA5 of the online appendix are devoted to precise statements.

1. A Group Bargaining Game

We consider a bargaining game related to those studied by Dávila et al. (2009) and Penta (2011). These papers are part of a broader literature that seeks multilateral bargaining foundations for Walrasian outcomes.²⁴

In the game, agents take turns speaking in sequence, and each agent can make a proposal about the ratios in which individuals should contribute. A typical proposal says, “For every unit done by me, I demand that agent 1 contribute 3 units, agent 2 contribute 0.5 units,” and so on. 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, a period of delay occurs and the next proposer gets to speak. Until an agreement is reached, players receive the payoff of the status quo outcome, and they discount at rate $\delta > 0$ per period.

The result is that the only Pareto efficient equilibrium outcomes involve immediate agreement on a centrality action profile. Thus, in a natural multilateral generalization of sequential bargaining, equilibrium play, along with the requirement of efficiency, selects the Lindahl outcome. The details are in Section OA3 of the online appendix.

2. Implementation Theory: The Lindahl Outcome as a Robust Selection

An alternative approach, based on implementation theory, places a more stringent normative requirement on the game—requiring all equilibria to yield efficient improvements on the status quo. It turns out that Lindahl outcomes play a distinguished role from this perspective as well. Again we sketch the result, relegating the formal treatment to Section OA4 of the online appendix.²⁵

A designer specifies a mechanism—message spaces for all the agents and an (enforceable) outcome function that maps messages into action profiles $\mathbf{a} \in \mathbb{R}^n$. The designer assumes that the profile of players’ preferences, $\mathbf{u} = (u_1, u_2, \dots, u_n)$, comes from some particular set \mathcal{U} , but she does not know exactly which preferences they will have in this set. She

²⁴ See also Yildiz (2003) and Dávila and Eeckhout (2008).

²⁵ Our result is analogous to theorem 3 of Hurwicz (1979a). Because the environment studied in that paper—with assumptions such as nonzero endowments of all private goods—is not readily adapted to our problem, we prove the result separately, using Hurwicz’s insights combined with Maskin’s theorem.

also assumes that players will end up playing a complete-information Nash equilibrium of her game, but she has no control over which equilibrium. We look for games the designer can create in which, for all preference profiles, all Nash equilibria satisfy Pareto efficiency and individual rationality. Pareto efficiency requires that any action profile resulting from equilibrium play of the game is Pareto efficient. Individual rationality ensures that every player is no worse off than at the status quo. We also require that equilibrium outcomes depend continuously on agents' preferences: arbitrarily small changes in preferences cannot force large changes in the equilibrium set.

It turns out that there are certain outcomes that occur as equilibrium outcomes for every mechanism satisfying the desiderata we have outlined above. This set of outcomes is called the set of *robustly attainable outcomes*. And the argument in favor of the Lindahl selection is the fact that, under suitable assumptions, this set is exactly equal to the set of Lindahl outcomes.

More precisely, the result is as follows. Assume that \mathcal{U} consists of all profiles satisfying the assumptions of Section II.B and that the number of players n is at least 3. Then the robustly attainable outcomes are the Lindahl outcomes.

To see why this result is useful, suppose that multiplicity of equilibria is considered a drawback of a mechanism—perhaps because this renders it less effective at coordinating the players on one efficient outcome. In that case, mechanisms implementing just the centrality action profiles do the best job of avoiding multiplicity. Such mechanisms exist exactly when there is a unique centrality action profile. In those cases, that is the outcome implemented.

3. Coalitional Deviations: A Core Property

As we are modeling negotiations, a natural question is whether some subset of the agents could do better by breaking up the negotiations and coming to some other agreement among themselves. Although this is outside the scope of actions available to the agents as modeled, the Lindahl outcomes are robust to coalitional deviations, if we assume that after a deviation, negotiations collapse and the nondeviating players choose their individually optimal responses, which are the status quo actions.²⁶ In our setting, this also minimizes the payoffs of any group of deviating players, taking the deviators' actions as given. Thus, the response by the complementary coalition is both individually optimal for the punishers and maximally harsh to the punished. We consider an outcome

²⁶ For any actions agents other than i can take, holding constant these actions \mathbf{a}_{-i} , agent i 's payoff is maximized by i selecting $a_i = 0$.

robust to coalitional deviations if no coalition would like to deviate, anticipating such a punishment.

Then we have the following result: if $\mathbf{a} \in \mathbb{R}_+^n$ has the centrality property, then \mathbf{a} is robust to coalitional deviations in the sense just described. This result is presented formally in Section OA5 of the online appendix

The remarkable yet simple argument for this, due to Shapley and Shubik (1969), is that the standard core of the artificial economy we presented above (with tradeable externalities) can be identified with the set of action profiles that are robust to coalitional deviations in our setting. In defining the core of the economy with tradeable externalities, we think of a deviating coalition ceasing trade with players outside of it. When externalities are not tradeable, we define outcomes robust to coalitional deviations by positing that a deviating coalition is punished by players outside the coalition reverting to the zero action level, that is, the action level at which the deviating coalition receives no benefits from the rest of society. Both coalitional deviations yield the same payoffs, so the same action profiles are robust to coalitional deviations in both settings.

4. A General Comment on Commitment and Information

The foundations for Lindahl outcomes that we have presented in this section have two key features: (1) commitment over actions and (2) complete information among the negotiating agents.

The assumption of commitment is standard in mechanism design and in our case is crucial for overcoming the free-riding problem. Some amount of commitment is necessary to contemplate efficient solutions—whether that commitment is credible because of the incentives created by repeated interaction or modeled via exogenous rules of the game, as in Sections IV.B.1 and IV.B.2. How much enforcement is possible in particular public goods problems is a critical question. Our contribution is to examine, in the benchmark case where there is commitment, how the network of externalities affects an important class of efficient solutions.

In terms of information, we assume that while the designer of the game may be ignorant of everything but the basic structure of the environment, the players interact with complete information about each other's preferences. Indeed, when transferable utility is not assumed—that is, when Vickrey-Clarke-Groves pivot mechanisms are not available—mechanism design with interim uncertainty in environments such as ours is not well understood.²⁷ Versions of our model with asymmetric information are cer-

²⁷ See, e.g., Garratt and Pycia (2016) for recent work.

tainly worth studying. We would expect the connections we identify between favor-trading games and networks to be relevant for that analysis.

V. Applications

In this section, we present four applications of our general results. First, we show how the above analysis can be used to predict who will be admitted to a team. Second, we extend the framework to endogenize the status quo, making it a Nash equilibrium of a unilateral-contribution game; this enriches the comparative statics of the problem, since now both the status quo and Lindahl outcomes move around with the environment. Third, we use special cases of our results to provide market interpretations of several measures of network centrality that have been utilized in a variety of settings, both within economics and especially in other fields. Finally, we study when a negotiating group or a team may be subdivided without much loss in terms of the outcome they reach.

For some of these discussions, it is helpful to think about eigenvector centralities in terms of walks on the network, so we begin with a discussion of that. In Section III.B, we saw that the spectral radius of the benefits matrix could be interpreted in terms of the values of long cycles. A related interpretation applies to centrality action profiles. A *walk of length ℓ in the matrix \mathbf{M}* is a sequence $(w(1), w(2), \dots, w(\ell + 1))$ of elements of N (agent indices) such that $M_{w(t)w(t+1)} > 0$ for each $t \in \{1, 2, \dots, \ell\}$.²⁸ Let $\mathcal{W}_i^\ell(\ell; \mathbf{M})$ be the set of all walks of length ℓ in \mathbf{M} ending at i (in our notation, such that $w(\ell + 1) = i$). For a nonnegative matrix \mathbf{M} , define the value of a walk w of length ℓ as the product of all matrix entries (i.e., link weights) along the walk:

$$v(w; \mathbf{M}) = \prod_{t=1}^{\ell} M_{w(t)w(t+1)}.$$

Note that such walks can repeat nodes—for example, they may cover the same cycle many times. Then we have the following.

PROPOSITION 3. Let $\mathbf{M} = \mathbf{B}(\mathbf{a})^\top$, and assume that this matrix is aperiodic.²⁹ Then \mathbf{a} has the centrality property if and only if, for every i and j ,

$$\frac{a_i}{a_j} = \lim_{\ell \rightarrow \infty} \frac{\sum_{w \in \mathcal{W}_i^\ell(\ell; \mathbf{M})} v(w; \mathbf{M})}{\sum_{w \in \mathcal{W}_j^\ell(\ell; \mathbf{M})} v(w; \mathbf{M})}.$$

A walk in $\mathbf{B}(\mathbf{a})^\top$ ending at i can be thought of as a chain of benefit flows: for example, k helps j , who helps i . The value of such a walk is the

²⁸ Note that a cycle is a special kind of walk.

²⁹ A *simple cycle* is one that has no repeated nodes except the initial/final one. A matrix is said to be *aperiodic* if the greatest common divisor of the lengths of all simple cycles in that matrix is 1.

product of the marginal benefits along its links. According to proposition 3, at a centrality action profile (and hence a Lindahl outcome), a player contributes in proportion to the total value of such benefit chains that end with him.³⁰

An implication of this analysis is that if the benefits i receives from j decrease at all action profiles—that is, $B_{ij}(\mathbf{a})$ decreases for all \mathbf{a} —then i 's centrality action level relative to that of all other agents will decrease. Thus, it is the benefits i receives, rather than the benefits i confers on others, that really matter for i 's Lindahl action. If, for example, there is an agent who can very efficiently provide benefits to the other agents and centrality action profiles are played, then there can be high returns from increasing the marginal benefits that this agent receives from others (and particularly those others with high eigenvector centrality). This has important implications, which we now discuss.

A. Application: Admitting a New Team Member

Suppose that agents $N = \{1, 2, 3\}$ currently constitute a team. These initial team members must decide whom, if anyone, to admit as a new member of their team. They have four options: admit nobody or admit a new team member $j \in M = \{4, 5, 6\}$. Afterward, the formed team collectively decides how much effort each of them should exert. We assume that these negotiations result in the Lindahl actions being played (see Sec. IV.B.1 for a motivation).

Who can provide benefits to whom in the initial team is described by the unweighted, directed graph \mathbf{G} (with entries in $\{0, 1\}$), illustrated in figure 3A. Once the decision about team composition has been made, G_{ij} is set to 0 if either i or j is not on the team. We assume that the original team members N can provide relatively strong benefits to each other and to the new team members M but that the new team members M are able to provide only weaker benefits. Specifically, the utility function of i is

$$u_i(\mathbf{a}) = \sum_{j \in N} G_{ij} \log(1 + a_j) + \sum_{j \in M} \frac{G_{ij}}{4} \log(1 + a_j) - a_i.$$

Agents not on the team will choose to exert no effort and will receive a payoff of 0.³¹ Figure 3B illustrates all possible benefit flows. Whom, if anyone, should the initial team members admit? Will the initial team members be able to agree on whom to admit?

³⁰ The formula of the proposition would also hold if had we defined $\mathbf{M} = \mathbf{B}(\mathbf{a})$ and replaced $\mathcal{W}_i^j(\ell; \mathbf{M})$ by $\mathcal{W}_i^j(\ell; \mathbf{M})$, which is the set of walks of length ℓ in \mathbf{M} that start at i . The convention we use above is in keeping with thinking of a walk in $\mathbf{B}(\mathbf{a})^\top$ capturing the direction in which benefits flow; recall the discussion in Sec. 3.2.

³¹ For such an agent i , $u_i(\mathbf{a}) = -a_i$.




FIG. 3.—*A*, The original negotiators and who among them can benefit whom. *B*, Benefits accruing to and coming from potential additional negotiators.

A quick inspection of figure 3 suggests that each original team member might most prefer admitting a new team member who can work with him directly. However, it is also worth noting that 3 is the only member of the original team that provides benefits to both of the other original team members. Increasing the effort of 3 could therefore be more efficient than increasing the effort of 1 or 2. Moreover, recall that as the Lindahl actions will be taken after the admission decision is made, those who receive higher marginal benefits will make more effort (by theorem 1 and the discussion at the start of this section). Perhaps, then, it might be relatively efficient to admit 6, the potential “helper” of 3, to induce 3 to take the highest possible action. It turns out that this is the case, and this increased efficiency exceeds the direct benefits 1 or 2 can receive from admitting 4 or 5. As a result, all the initial team members’ interests are actually aligned.

We now formalize this intuition, using the tools we have developed. By theorem 1, agents’ Lindahl actions are given by their centralities in **B(a)**. Applying the scaling-indifference characterization of these actions, $\mathbf{J}(\mathbf{a})\mathbf{a} = \mathbf{0}$, we find that the centrality action of agent i is characterized by $a_i = \sum_j (G_{ij}a_j)/(1 + a_j)$. The unique³² centrality actions if no new team members are admitted are \mathbf{a} ; if 4 is admitted they are \mathbf{a}' ; if 5 is instead added, they are \mathbf{a}'' ; and if 6 is added, they are \mathbf{a}''' , where in each vector the last entry corresponds to the action taken by the new team member:

$$\mathbf{a} = \begin{pmatrix} .408 \\ .225 \\ .290 \\ — \end{pmatrix} \quad \mathbf{a}' = \begin{pmatrix} .523 \\ .256 \\ .343 \\ .343 \end{pmatrix} \quad \mathbf{a}'' = \begin{pmatrix} .462 \\ .286 \\ .316 \\ .222 \end{pmatrix} \quad \mathbf{a}''' = \begin{pmatrix} .497 \\ .279 \\ .386 \\ .279 \end{pmatrix}.$$

If added, 4 will take a higher action than 6, who will take a higher action than 5. However, the inclusion of 6 induces agent 3 to take the high-

³² Uniqueness is established by noting that each agent’s preferences satisfy the gross substitutes property, and therefore the Lindahl outcome is unique (McKenzie 1960).

est action, providing indirect benefits to both 1 and 2. The utility vectors for the original negotiators, when the centrality action profiles are played, are shown below for the options of, in order, admitting nobody, admitting 4, admitting 5, and admitting 6:

$$\begin{aligned} \mathbf{u}(\mathbf{a}) &= \begin{pmatrix} .049 \\ .030 \\ .052 \end{pmatrix} & \mathbf{u}(\mathbf{a}') &= \begin{pmatrix} .074 \\ .040 \\ .077 \end{pmatrix} & \mathbf{u}(\mathbf{a}'') &= \begin{pmatrix} .064 \\ .039 \\ .064 \end{pmatrix} \\ \mathbf{u}(\mathbf{a}''') &= \begin{pmatrix} .076 \\ .048 \\ .078 \end{pmatrix}. \end{aligned}$$

Thus, the incentives of the original team members are perfectly aligned. Even though different potential additions benefit different original team members, all prefer admitting 6 to admitting 4 to admitting 5 to admitting nobody. The indirect benefit flows from admitting agent 6 outweigh the direct benefit flows agents 1 and 2 would receive from admitting agent 4 or 5.

While in general the incentives of agents need not be aligned when deciding whom to include in a team, studying the network structure of the externalities can help us understand the implications of including different team members. One general lesson is that team members who have the potential to provide benefits to many others realize this potential when they are the beneficiaries of links from new members.

B. Endogenous Status Quo

In the main analysis, we made assumptions so that the Nash equilibrium action profile was $\mathbf{0}$ (the corner of the nonnegative orthant, in which actions lie) and argued that this was essentially a normalization. While convenient for some purposes, this setup is not suited for others, such as studying the difference between Nash and Lindahl outcomes when changes in the environment cause both to change. In this section, we endogenize the status quo, making it the Nash equilibrium of a simultaneous-move public goods contribution game. To this end, suppose that assumptions 2 (positive externalities), 3 (irreducibility), and 4 (bounded improvements) continue to hold but that we relax assumption 1 (costly actions). Consider now a simultaneous-move game in which each agent chooses an action $a_i \in \mathbb{R}_+$.

Consider a Nash equilibrium action profile \mathbf{a}^{NE} , defined by the condition that $a_i^{\text{NE}} = \text{argmax}_{a_i} u_i(a_i, \mathbf{a}_{-i}^{\text{NE}})$. By the concavity of the utility functions, for all i , we have $\partial u_i(\mathbf{a}^{\text{NE}})/\partial a_i \leq 0$, with equality holding if $a_i^{\text{NE}} > 0$.

Take the actions \mathbf{a}^{NE} as the status quo. For any \mathbf{a} , define $\hat{\mathbf{a}} := \mathbf{a} - \mathbf{a}^{\text{NE}}$ to be the increment of a given action profile over the Nash equilibrium action profile. Also, define the utility profile $\hat{\mathbf{u}}(\hat{\mathbf{a}}) := \mathbf{u}(\mathbf{a}^{\text{NE}} + \hat{\mathbf{a}})$. With the action space and utility profiles reparameterized in this way, the status quo action profile $\hat{\mathbf{a}} = \mathbf{0}$ is the Nash equilibrium action profile, and the ratio $\partial \hat{u}_i(\mathbf{0}) / \partial \hat{a}_i \leq 0$ for all i . The assumption $\hat{\mathbf{a}} \geq \mathbf{0}$, which is maintained for our analysis, entails that players do not take actions below their Nash equilibrium actions.³³

There are then two cases to consider. If $\partial \hat{u}_i(\mathbf{0}) / \partial \hat{a}_i < 0$ for all i (case I), then assumptions 1–4 all hold for the environment given by $\hat{\mathbf{u}}$, and our results go through unchanged. If $\partial \hat{u}_i(\mathbf{0}) / \partial \hat{a}_i = 0$ for some i (case II), then assumption 1 will be violated. This prevents us from directly applying our results. However, this is a technical rather than a substantive problem, as we show now. Indeed, in case II, the proofs of proposition 1 and theorem 1 go through with some modification.

For proposition 1, to show that Pareto efficiency of an interior \mathbf{a}^* implies that $\mathbf{B}(\mathbf{a}^*; \hat{\mathbf{u}})$ has spectral radius 1, we can again start by looking at the first-order condition $\theta \mathbf{J}(\mathbf{a}^*; \hat{\mathbf{u}}) = \mathbf{0}$ for some nonzero $\theta \in \mathbb{R}_+^n$. We cannot immediately divide each row of this equation by $-J_{ii}(\mathbf{a}^*; \hat{\mathbf{u}}) = -\partial \hat{u}_i(\mathbf{a}^*; \hat{\mathbf{u}}) / \partial a_i$ to convert this into $\theta \mathbf{B}(\mathbf{a}^*; \hat{\mathbf{u}}) = \theta$, because that might involve dividing by 0. So first we argue that any \mathbf{a}^* satisfying the first-order conditions must have $J_{ii}(\mathbf{a}^*; \hat{\mathbf{u}}) < 0$ for every i . This is precisely the content of the following lemma, whose proof is deferred to appendix C.

LEMMA 1. Take any utility profile $\hat{\mathbf{u}}$ satisfying assumptions 2 and 3, with $\partial \hat{u}_i(\hat{\mathbf{a}}; \hat{\mathbf{u}}) / \partial \hat{a}_i \leq 0$ for every i and every $\hat{\mathbf{a}}$. If the first-order condition $\theta \mathbf{J}(\mathbf{a}^*; \hat{\mathbf{u}}) = \mathbf{0}$ holds for a nonzero vector of Pareto weights, $\theta \in \mathbb{R}_+^n$, then $J_{ii}(\mathbf{a}^*; \hat{\mathbf{u}}) < 0$ for every i .

With this lemma in hand, the proof of proposition 1 can continue as before. The intuition for the lemma is simple: for any $\hat{\mathbf{a}}$, we can construct a new $\tilde{\mathbf{u}}$ so that $J_{ii}(\hat{\mathbf{a}}; \tilde{\mathbf{u}})$ is negative but very small whenever it was zero under $\hat{\mathbf{u}}$, and $\mathbf{J}(\hat{\mathbf{a}}; \tilde{\mathbf{u}})$ is unchanged otherwise. Now $\mathbf{B}(\hat{\mathbf{a}}; \tilde{\mathbf{u}})$ is irreducible and thus contains cycles; by making the cost we have introduced sufficiently small, we can make the value of these cycles very large, which, as shown in Section III.B, creates a Pareto improvement, guaranteeing that $\hat{\mathbf{a}}$ is not an efficient point under $\tilde{\mathbf{u}}$. It is therefore not efficient under $\hat{\mathbf{u}}$, either.

The following corollary shows that nonzero Nash equilibrium action profiles are inefficient.

COROLLARY 1. The Nash equilibrium action profile is Pareto efficient only if it is the zero action profile.

³³ This restriction makes sense when investments are irreversible, though our main results would have analogues without it.

Proof. Suppose that there is a Nash equilibrium action profile (before parameterization) \mathbf{a}^{NE} so that $a_i^{\text{NE}} > 0$ for some i . Then $J_{ii}(\mathbf{a}^{\text{NE}}; \mathbf{u}) = J_{ii}(\mathbf{0}; \hat{\mathbf{u}}) = 0$. So by lemma 1, \mathbf{a}^{NE} cannot solve the Pareto problem. QED

When the Nash equilibrium action profile is zero, every agent might be up against the lower bound and might prefer to take lower actions that are unavailable. This could occur if actions correspond to irreversible investments (perhaps sunk investments in clean energy), a Nash equilibrium is played, and then some parameter of the environment changes, causing the Nash equilibrium actions to decrease. Then, even though the positive benefits of higher actions are neglected in the private decision of how much effort to exert, the zero action can be Pareto efficient. The positive marginal benefits that would accrue to others may not cover the private marginal costs of higher actions.

Like proposition 1, theorem 1 extends to the case in which the status quo actions are a Nash equilibrium. Indeed, the existing proofs go through, using the newly strengthened proposition 1 we have just discussed.

Having handled the technical issues in defining our solution with a general Nash equilibrium status quo, we can draw on the network games literature pioneered by Ballester et al. (2006) to characterize the Nash equilibrium status quo and compare it to the Lindahl action profile. However, to use results from this literature, it is convenient to look at a special case of our setting. Let \mathbf{G} be an undirected, unweighted graph ($g_{ij} = g_{ji} \in \{0, 1\}$) with no self-links ($g_{ii} = 0$), describing which agents are neighbors. For a matrix \mathbf{M} we let

$$\lambda_{\min}(\mathbf{M}) := \min\{|\lambda| : \lambda \text{ is an eigenvalue of } \mathbf{M}\}.$$

Suppose that utility functions are given by

$$u_i(\mathbf{a}) = b\left(a_i + \delta \sum_j g_{ij} a_j\right) - a_i,$$

where b is a strictly increasing and convex function and $\lambda_{\min}(\mathbf{G}) < 1/\delta$. Note that in this formulation, an agent's neighbors' actions are perfect substitutes for one another.

PROPOSITION 4 (Bramoullé et al. 2014). Given the above assumptions, there is a unique Nash equilibrium.

To further compare the Nash equilibrium and Lindahl outcomes, let \mathbf{G} be a regular graph in which all agents have k links (and continue to assume that $\lambda_{\min}(\mathbf{G}) < 1/\delta$). Let $\beta(x) = b'(x)$. The unique Nash equilibrium is then symmetric, with $a^{\text{NE}} = a_i^{\text{NE}} = \beta^{-1}(1)/(1 + \delta k)$. Given the above utilities, at an action profile where the benefits matrix is well defined, for $i \neq j$ the ij th entry of the benefits matrix is

$$B_{ij}(\mathbf{a}) = g_{ij} \frac{\delta\beta(a_i + \delta\sum_j g_{ij}a_j)}{[\beta(a_i + \delta\sum_j g_{ij}a_j) - 1]}.$$

We now apply theorem 1. It is helpful to write Lindahl equilibrium actions as $\mathbf{a}^{\text{LE}} = \mathbf{a}^{\text{NE}} + \hat{\mathbf{a}}$. As the Nash equilibrium is the status quo, theorem 1 then adjusts so that $\mathbf{B}(\mathbf{a}^{\text{LE}})\hat{\mathbf{a}} = \hat{\mathbf{a}}$ or, equivalently, $\mathbf{B}(\mathbf{a}^{\text{LE}})(\mathbf{a}^{\text{LE}} - \mathbf{a}^{\text{NE}}) = (\mathbf{a}^{\text{LE}} - \mathbf{a}^{\text{NE}})$. We therefore have that

$$a_i^{\text{LE}} - a_i^{\text{NE}} = \sum_j g_{ij} \frac{\delta\beta(a_i^{\text{LE}} + \delta\sum_j g_{ij}a_j^{\text{LE}})}{1 - \beta(a_i^{\text{LE}} + \delta\sum_j g_{ij}a_j^{\text{LE}})} (a_j^{\text{LE}} - a_j^{\text{NE}}),$$

for all i . There is then a Lindahl equilibrium in which all agents take the same action and³⁴

$$a^{\text{LE}} = a_i^{\text{LE}} = \frac{\beta^{-1}(1/(1+k\delta))}{1+k\delta} > \frac{\beta^{-1}(1)}{1+\delta k} = a_i^{\text{NE}} = a^{\text{NE}}.$$

The Nash equilibrium actions are decreasing in k and δ . Intuitively, at higher values of k and δ there is more free riding on the actions of other agents, and agents choose lower actions in equilibrium. Interestingly, even though the increased externalities are internalized in the Lindahl equilibrium, the Lindahl actions are not always increasing in the level of externalities δ .³⁵ For example, when $b(\cdot) = \log(\cdot)$, the Lindahl actions are invariant with respect to δ (and k). The intuition underlying this is that when the benefits function is sufficiently concave, increasing δ at the current Lindahl equilibrium increases the agents' consumption of the public goods, causing the marginal benefits from further consumption to decrease sufficiently that in the Lindahl equilibrium agents reduce their actions. Nevertheless, the presence of free riding in the Nash equilibrium in comparison to the Lindahl equilibrium is observable when comparing the Lindahl and Nash actions. The ratio of the Lindahl to the Nash actions is increasing in k , δ , and the concavity of the b .³⁶

³⁴ This calculation relies heavily on the symmetry of this problem. It would be interesting to explore the difference between Lindahl and Nash outcomes more generally, although also harder because the key quantities will only be implicitly defined.

³⁵ Implicitly differentiating the first-order condition and rearranging gives

$$\frac{\partial a_i^{\text{LE}}}{\partial \delta} = -\frac{k}{b''(a(1+k\delta))(1+k\delta)^3} [1 + ab''(a(1+k\delta))(1+k\delta)^2], \text{ where } \mathbf{a} = \mathbf{a}^{\text{LE}}.$$

This expression is weakly greater than zero if and only if $-b''(a(1+k\delta)) \leq 1/[a(1+k\delta)^2]$.

³⁶ As b is concave, β is strictly decreasing, so β^{-1} is well defined and also strictly decreasing. It follows that $\beta^{-1}(1/(1+k\delta))$ is increasing in k and δ .

It is clear from the calculations that regularity played an essential role in this example. To examine the same questions without imposing regularity, in Section OA7 of the online appendix, we consider a star graph, such that one center agent is connected to all other agents, who have only this link. Setting $b(\cdot) = \log(\cdot)$, a specification under which any regular graph has Lindahl actions invariant in the level of externalities δ , we find that the center agent's Lindahl action decreases in δ , while the periphery agents' Lindahl actions are increasing in δ . The ratio of the Lindahl to Nash actions, as with regular graphs, is increasing in δ .

C. Explicit Formulas for Lindahl Outcomes

Several measures of network centrality have been extensively employed in the networks literature. In this section we use our results to provide new foundations for three of them. We do so by linking each measure to the Lindahl equilibrium under different parametric assumptions on preferences.

The preferences we consider are

$$u_i(\mathbf{a}) = -a_i + \sum_j (\alpha G_{ij} a_j + H_{ij} \log a_j), \quad (6)$$

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 ($\mathbf{a} = \mathbf{B}(\mathbf{a})\mathbf{a}$) discussed throughout the paper boils down to $\mathbf{a} = \mathbf{h} + \alpha \mathbf{G}\mathbf{a}$.

Several special cases are worth considering. If $\alpha = 0$, then $a_i = h_i$, and i 's Lindahl action is equal to the total weight of i 's incoming links in \mathbf{H} . This measure of i 's centrality in the network \mathbf{H} is known as i 's *degree centrality*. If, instead, $h_i = 1$ for all i , then agents' Lindahl actions are $\mathbf{a} = (\mathbf{I} - \alpha \mathbf{G})^{-1} \mathbf{1}$. The right-hand side is a different measure of agents' centralities in the network \mathbf{G} , known as their *Bonacich centralities*. Like degree centrality, Bonacich centrality depends on the number of i 's neighbors, but it also depends on longer-range paths.³⁷ Finally, in this setting, as α approaches 1, agents' actions become proportional to their normalized eigenvector centralities in \mathbf{G} . These results are further discussed in Section OA8 of the online appendix.

As Lindahl outcomes are defined in terms of prices, the formulas we have presented may be viewed as microfoundations or interpretations of

³⁷ It can also be characterized via the equation $\mathbf{a} = \alpha \mathbf{G}\mathbf{a} + \mathbf{1}$, which resembles the condition defining eigenvector centrality. For more background and discussion, see Ballester et al. (2006, sec. 3) and Jackson (2008, sec. 2.2.4).

network centrality measures in terms of price equilibria. Each result says that for particular preferences, the allocations defined by Lindahl are equal to centralities, according to a corresponding measure. Such a connection permits new interpretations of well-known centrality measures. Beyond these interpretations, the connection between centrality measures and prices may permit new analytical techniques inspired by price equilibria.

*D. Approximating the Full Benefits of Negotiation
with Smaller Groups*

There are often large costs involved in organizing a large multilateral negotiation, and therefore it is important to know when most of the benefits of negotiating can be achieved by instead organizing negotiations in smaller groups. Our framework allows us to give a simple analysis of the costs of subdividing a negotiation.³⁸

We consider an arbitrary Pareto efficient outcome \mathbf{a}^* that a planner would like to achieve. We then suppose that the agents are divided into two subsets, M and M^c , and that \mathbf{a}^* is proposed to each. Then, each group can contemplate deviations from \mathbf{a}^* that are Pareto improving for that group. A group will generally have a Pareto-improving deviation of reducing efforts relative to \mathbf{a}^* , because as a group they pay all the costs of effort but do not internalize any of the benefits to the complement.³⁹

How cheaply can a planner incentivize agents to stay with the original outcome rather than deviate? To quantify the cost of such incentives, we imagine that the social planner can subsidize individuals' effort, and we ask when only a small amount of subsidy will be required to remove any incentive for each group to move away from the target efficient point \mathbf{a}^* .

To that end, we set $J_{ii}(\mathbf{a}) = -1$ for each i and all \mathbf{a} , and we moreover assume that there is a numeraire in which each agent could be paid—one that enters his utility additively. We do not allow transfers among the agents, but we do allow a planner to use transfers of this numeraire (potentially required to be "small" in some sense) to subsidize individuals' efforts. Thus, we posit that the planner can modify the environment to one with payoff functions,

$$\tilde{u}_i(\mathbf{a}) = u_i(\mathbf{a}) + m_i(\mathbf{a}),$$

³⁸ We are grateful to an anonymous reviewer for suggesting this analysis.

³⁹ To prove this formally, one can use a strict version of fact 1b to show that the benefits matrix restricted to just one group has a largest eigenvalue strictly less than 1 (assuming that the benefits matrix among the grand coalition was irreducible) and then use proposition 6 (app. B) to show that some reduction of all actions yields a Pareto improvement.

where $m_i(\mathbf{a})$ must be nonnegative. We say that 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 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 5. Consider a Pareto efficient outcome \mathbf{a}^* , and let θ be the corresponding Pareto weights. Then,

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

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 .

In graph theory terms, this is the weight of the cut M in a weighted graph derived from $\mathbf{B}(\mathbf{a}^*)$, whose edge weights are $W_{ij} = (\theta_i/\theta_j) B_{ij}(\mathbf{a}^*) a_j^*$. When \mathbf{a}^* and θ are held fixed, the bound in the proposition becomes small if the network given by $\mathbf{B}(\mathbf{a}^*)$ has only a small total weight on links across groups. Note that it is the properties of marginal benefits that are key—given this result, 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. One conclusion is that if there is another 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).⁴⁰ (The difference between the largest and second-largest eigenvalues is often referred to as the *spectral gap*.) Thus, eigenvalues of $\mathbf{B}(\mathbf{a}^*)$ other than the largest have economic implications in our setting.

VI. Concluding Discussion

In this section, we discuss the extent to which some of our more economically restrictive assumptions can be relaxed, elaborate on how our work fits into several related literatures, and offer some concluding remarks.

A. Relaxing Assumptions

The assumption of a single dimension of effort per agent is relaxed in Section OAI of the online appendix, which introduces a benefits matrix for each dimension and characterizes efficient outcomes via the eigen-

⁴⁰ For a survey of some related results, see von Luxburg (2007).

vectors and eigenvalues of these matrices and Lindahl outcomes via scaling indifference. The implicit assumption of no transfers of a numeraire (“side payments” separate from the actions) is relaxed in Section OA2 of the online appendix, where we give the analog of the Samuelson condition from public finance in our setting.

An important and restrictive assumption we make is that all externalities are positive. This environment is equivalent to one with negative externalities in which it is costly to decrease actions. For example, in our simple example in Section I, the action towns take can be seen as reducing their pollution by producing less.

The case of both positive and negative externalities is more challenging, and we now discuss the extent to which assumption 2 can be relaxed. The key mathematical result we lean on throughout our analysis is the Perron-Frobenius theorem, which applies only to nonnegative matrices. However, there are generalizations of the theorem in which the assumption of nonnegativity is weakened (see, e.g., Johnson and Tarazaga 2004; Noutsos 2006). The weaker assumptions essentially require that the positive externalities dominate the negative externalities. For example, one sufficient condition is that all entries of $\mathbf{B}'(\mathbf{a})$ are positive for all sufficiently large ℓ , which is related to walks in the network (see Sec. V). We consider the more restrictive environment only for simplicity.

B. Related Literature

A recent literature has found a connection between the Nash equilibria of one-shot games in networks and centrality measures in those networks. Key papers in this literature include Ballester et al. (2006) on skill investment with externalities and Bramoullé et al. (2014) on local public goods. Most recently, Allouch (2015) has studied a network version of the setting introduced by Bergstrom, Blume, and Varian (1986) on the voluntary (static Nash) private provision of public goods. Generalizing results of Bramoullé et al. (2014), he derives comparative statics of public goods provision using network centrality tools.⁴¹ Unlike our approach, results in this literature typically require best responses to take a particular parametric form.⁴² Another, more fundamental difference is that the games we focus on in Section IV.B.1 are designed to overcome the free riding present in the private provision models; in contrast, the games studied in the papers mentioned above do not share this feature, and so the Nash equilibria are typically inefficient (see Sec. V.B for more on this). A recent paper

⁴¹ These papers contain more complete discussions of this literature. See also Bramoullé and Kranton (2007).

⁴² For some recent work in which parametric assumptions have been relaxed in the context of network formation, see Baetz (2015) and Hiller (2017).

from that literature, perhaps closest to our work insofar as network centrality is related to prices in a market, is Chen, Zenou, and Zhou (2018). There, two firms each offer a different substitutable product to consumers embedded in a network where consumers' utilities depend on their neighbors' consumptions. The firms can price discriminate, and using the technology developed by Ballester et al. (2006), equilibrium prices in this market are tied to agents' centralities in the network. Key differences remain insofar as the markets in that paper are not competitive and decisions are unilateral and only privately optimal.

In emphasizing the correspondence between centrality and outcomes of a competitive market, our perspective is related to that of Du et al. (2015), who find a connection between eigenvector centrality and outcomes in an exchange economy with Cobb-Douglas preferences. The parametric forms required to relate outcomes explicitly to familiar centrality measures differ in the two models, but both papers share the perspective that centrality and markets are closely related, and each concept can be used to shed light on the other. An advantage of the public goods economy we study is that our explicit characterizations discussed in Section V.C are a special case of an eigenvector characterization that applies without parametric assumptions. We believe that these projects, taken together, offer hope for a fairly rich theory of connections between market outcomes and network centrality.

Conitzer and Sandholm (2004) study "charity auctions," which, like the strategic settings we discuss in Section IV.B, are intended to implement Pareto improvements in the presence of externalities. In that model, agents condition their charitable contributions on others' contributions and so choose action vectors that are reminiscent of the directions chosen in the bargaining game of Section IV.B.1. A paper taking this approach in a network context is Ghosh and Mahdian (2008). Their model locates people on a social network and assumes that they benefit linearly from their neighbors' contributions, with a cap on how much any individual can contribute. There is an equilibrium of their game that achieves the maximum possible feasible contributions (subject to individual rationality), and this involves positive contributions being made if and only if the largest eigenvalue of the fixed network is greater than 1.

Understanding how the presence of externalities affects classical solutions (often ones inspired by markets) is an active area of research more broadly. For instance, a recent contribution by Pycia and Yenmez (2015) generalizes classical matching algorithms and characterizations to settings with both positive and negative externalities.

C. Conclusions

Many practical problems, such as preventing harmful climate change, entail a tragedy of the commons. It is in each agent's interest to free-ride

on the efforts of others. A question at the heart of economics, and of intense public interest, is the extent to which negotiations can overcome such problems and lead to outcomes different from, and better than, the outcomes under noncooperative, voluntary provision. Our thesis is that, in addressing this problem, it is informative to study the properties of a network of externalities.

Cycles in this network are necessary for there to be any scope for a Pareto improvement, and summing these cycles in a certain way identifies whether a Pareto improvement is possible or not. We can use this insight to identify which agents, or sets of agents, are essential to a negotiation in the sense that their participation is necessary for achieving a Pareto improvement on the status quo.

Moreover, a measure of how central agents are in this network—eigenvector centrality—characterizes which actions agents would take under the Lindahl solution. In our environment, the Lindahl solution is more than just a hypothetical construct describing what we could expect if missing markets were somehow completed. The Lindahl outcomes correspond to the efficient equilibria of a bargaining game. Moreover, an implementation-theoretic analysis selects the Lindahl solutions as ones that are particularly robust to the specification of the negotiation game.

From the eigenvector centrality characterization of Lindahl outcomes, we can see that agents' actions are determined by a weighted sum of the marginal benefits they receive, as opposed to the marginal benefits they can provide to others.⁴³ This has implications for the design of negotiations. If there is an agent who is in a particularly strong position to provide direct and indirect benefits to others, it will be especially important to include others in the negotiation who can help this agent. Our results formalize this intuition and can be used to quantify the associated trade-offs in the formation of a team.

Several interesting questions remain unanswered. Our focus on efficient outcomes presupposes group cooperation; but if the group can cooperate and commit as a whole, it is worth worrying about the possibility that a subset of the agents may coordinate on a deviation from some desired outcome. In Section IV.B.3, we note that Lindahl outcomes are robust to coalitional deviations, assuming that nondeviators revert to no effort, but realistic consequences of a coalition's renegeing are more complex. Are there efficient outcomes that are robust to deviations even with a richer model of the postrenegeing subgame, and how do such outcomes relate to properties of the benefits network?

Other natural questions include What incentives are there for investments that increase the benefits agents can confer on each other? In

⁴³ The most explicit version of this statement is in Sec. OA8.3 of the online appendix, in which we calculate that the weights of incoming walks according to an exogenous network fully determine equilibrium efforts.

what sense does the spectral radius provide an appropriate measure of how much scope for cooperation there is? In applications such as trade liberalization, where there are multiple actions available to the different agents, how should negotiations be designed?

Appendix A

The Perron-Frobenius Theorem

The key mathematical tool we use is the Perron-Frobenius theorem. We state it here for ease of reference and to enumerate the various parts of it that we rely on at different points in the paper.⁴⁴

THEOREM (Perron-Frobenius). Let \mathbf{M} be an irreducible, square matrix with no negative entries and spectral radius $r(\mathbf{M})$. Then,

- (1) The real number $r(\mathbf{M})$ is an eigenvalue of \mathbf{M} .
- (2) There is a vector \mathbf{p} (called a Perron vector) with only positive entries such that $\mathbf{M}\mathbf{p} = r(\mathbf{M})\mathbf{p}$.
- (3) If \mathbf{v} is a nonzero vector with nonnegative entries such that $\mathbf{M}\mathbf{v} = q\mathbf{v}$ for some $q \in \mathbb{R}$, then \mathbf{v} is a positive scalar multiple of \mathbf{p} , and $q = r(\mathbf{M})$.

Note that because a matrix has exactly the same eigenvalues as its transpose, all the same statements are true, with the same eigenvalue $r(\mathbf{M}) = r(\mathbf{M}^\top)$, when we replace \mathbf{M} by its transpose \mathbf{M}^\top . This observation yields a left-hand Perron eigenvector of \mathbf{M} , that is, a row vector \mathbf{w} such that $\mathbf{w}\mathbf{M} = r(\mathbf{M})\mathbf{w}$. For nonsymmetric matrices, it is typically the case that $\mathbf{w}^\top \neq \mathbf{p}$. The analogue of statement 3 in the theorem holds for \mathbf{w} .

Appendix B

Egalitarian Pareto Improvements

This appendix serves two purposes. First, it presents a result that is of interest in its own right, clarifying the sense in which the spectral radius of the benefits matrix measures the magnitude of inefficiency rather than merely diagnosing that there is some inefficiency. Second, it introduces some terminology and results that are useful in subsequent proofs, particularly the proof of proposition 2, which establishes the existence of a centrality action profile.

Let Δ_n denote the simplex in \mathbb{R}_+^n defined by $\Delta_n = \{\mathbf{d} \in \mathbb{R}_+^n : \sum_i d_i = 1\}$.

DEFINITION 4. The *bang-for-the-buck* vector $\mathbf{b}(\mathbf{a}, \mathbf{d})$ at an action profile \mathbf{a} along a direction $\mathbf{d} \in \Delta_n$ is defined by

$$b_i(\mathbf{a}, \mathbf{d}) = \frac{\sum_{j: j \neq i} J_{ij}(\mathbf{a}) d_j}{-J_{ii}(\mathbf{a}) d_i}.$$

⁴⁴ Meyer (2000, sec. 8.3) has a comprehensive exposition of this theorem, its proof, and related results.

This is the ratio

$$\frac{i\text{'s marginal benefit}}{i\text{'s marginal cost}}$$

evaluated at \mathbf{a} , when everyone increases actions slightly in the direction \mathbf{d} . We say that a direction $\mathbf{d} \in \Delta_n$ is *egalitarian at \mathbf{a}* if all the entries of $b_i(\mathbf{a}, \mathbf{d})$ are equal.

PROPOSITION 6. At any \mathbf{a} , there is a unique egalitarian direction $\mathbf{d}^{\text{eg}}(\mathbf{a})$. Every entry of $\mathbf{b}(\mathbf{a}, \mathbf{d}^{\text{eg}}(\mathbf{a}))$ is equal to the spectral radius of $\mathbf{B}(\mathbf{a})$.

Proposition 6 shows that for any action profile \mathbf{a} , there is a unique egalitarian direction in which actions can be changed at \mathbf{a} to equalize the marginal benefits per unit of marginal cost accruing to each agent and that this benefit-to-cost ratio will be equal to the spectral radius of $\mathbf{B}(\mathbf{a})$. Thus, the spectral radius of $\mathbf{B}(\mathbf{a})$, when it exceeds 1, is a measure of the size of egalitarian Pareto improvements available by increasing actions. (A corresponding interpretation applies when the spectral radius is less than 1.)

Proof. Fix \mathbf{a} and denote by r the spectral radius of $\mathbf{B}(\mathbf{a})$. Since $\mathbf{B}(\mathbf{a})$ is nonnegative and irreducible, the Perron-Frobenius theorem guarantees that $\mathbf{B}(\mathbf{a})$ has a right-hand eigenvector \mathbf{d} such that

$$\mathbf{B}(\mathbf{a})\mathbf{d} = r\mathbf{d}. \quad (\text{B1})$$

This is equivalent to $\mathbf{b}(\mathbf{a}, \mathbf{d}) = r\mathbf{1}$, where $\mathbf{1}$ is the column vector of ones. Therefore, there is an egalitarian direction that generates a bang for the buck of r (the spectral radius of $\mathbf{B}(\mathbf{a})$) for everyone.

Now suppose that $\tilde{\mathbf{d}} \in \Delta_n$ is any egalitarian direction; that is, for some b we have

$$\mathbf{b}(\mathbf{a}, \tilde{\mathbf{d}}) = b\mathbf{1}.$$

This implies

$$\mathbf{B}(\mathbf{a})\tilde{\mathbf{d}} = b\tilde{\mathbf{d}}. \quad (\text{B2})$$

By the Perron-Frobenius theorem (statement 3), the only real number b and vector $\tilde{\mathbf{d}} \in \Delta_n$ satisfying equation (B2) are $b = r$ and $\tilde{\mathbf{d}} = \mathbf{d}$.

Thus, $\mathbf{d}^{\text{eg}}(\mathbf{a}) = \mathbf{d}$ has all the properties claimed in the proposition's statement. QED

Appendix C

Omitted Proofs

Proof of Proposition 1

We first prove part a. For any nonzero $\boldsymbol{\theta} \in \mathbb{R}_+^n$, define $\mathcal{P}(\boldsymbol{\theta})$, the Pareto problem with Pareto weights $\boldsymbol{\theta}$, as

$$\text{maximize } \sum_{i \in N} \theta_i u_i(\mathbf{a}) \text{ subject to } \mathbf{a} \in \mathbb{R}_+^n.$$

Suppose that an interior action profile \mathbf{a}^* is Pareto efficient. Assumption 1 guarantees that $J_{ii}(\mathbf{a}^*)$ is strictly negative. We may multiply utility functions by positive

constants to achieve the normalization $J_{ii}(\mathbf{a}^*) = -1$ for each i . This is without loss of generality. It clearly does not affect Pareto efficiency, and it is easy to see that scaling utility functions does not affect $\mathbf{B}(\mathbf{a}^*)$. Since \mathbf{a}^* is Pareto efficient, it solves $\mathcal{P}(\theta)$ for some nonzero $\theta \in \mathbb{R}_+^n$ (this is a standard fact for concave problems). And therefore \mathbf{a}^* satisfies $\mathcal{P}(\theta)$'s system of first-order conditions: $\theta \mathbf{J}(\mathbf{a}^*) = \mathbf{0}$. By our normalization, $\mathbf{J}(\mathbf{a}) = \mathbf{B}(\mathbf{a}) - \mathbf{I}$, where \mathbf{I} is the n -by- n identity matrix, so the system of first-order conditions is equivalent to $\theta \mathbf{B}(\mathbf{a}^*) = \theta$.

This equation says that $\mathbf{B}(\mathbf{a}^*)$ has an eigenvalue of 1 with corresponding left-hand eigenvector θ . Since $\mathbf{B}(\mathbf{a}^*)$ is a nonnegative matrix, and is irreducible by assumption 3, the Perron-Frobenius theorem applies to it. That theorem says that the only eigenvalue of $\mathbf{B}(\mathbf{a}^*)$ that can be associated with the nonnegative eigenvector θ is the spectral radius itself.⁴⁵ Thus, the spectral radius of $\mathbf{B}(\mathbf{a}^*)$ must be 1.

Conversely, suppose that $\mathbf{B}(\mathbf{a}^*)$ has a spectral radius of 1, and again normalize each i 's utility function so that $J_{ii}(\mathbf{a}^*) = -1$. The Perron-Frobenius theorem guarantees that $\mathbf{B}(\mathbf{a}^*)$ has 1 as an eigenvalue and also yields the existence of a nonnegative left-hand eigenvector θ such that $\theta \mathbf{B}(\mathbf{a}^*) = \theta$. Consequently, the first-order conditions of the Pareto problem $\mathcal{P}(\theta)$ are satisfied (using the manipulation of the first-order conditions we used above). By the assumption of concave utilities, it follows that \mathbf{a}^* solves the Pareto problem for weights θ (i.e., the first-order conditions are sufficient for optimality), and so \mathbf{a}^* is Pareto efficient.

We now prove part b, starting with the case in which assumption 3 holds.

If $r(\mathbf{B}(\mathbf{0})) > 1$, then proposition 6 in appendix B yields an egalitarian direction at $\mathbf{0}$ with bang for the buck exceeding 1; this is a Pareto improvement at $\mathbf{0}$.

If $\mathbf{0}$ is not Pareto efficient, there is an $\mathbf{a}' \in \mathbb{R}_+^n$ such that $u_i(\mathbf{a}') \geq u_i(\mathbf{0})$ for each i , with strict inequality for some i . Using assumption 3, namely, the irreducibility of $\mathbf{B}(\mathbf{a}')$, as well as the continuity of the u_i , we can find an \mathbf{a}'' with all positive entries so that $u_i(\mathbf{a}'') > u_i(\mathbf{0})$ for all i .⁴⁶ Let \mathbf{v} denote the derivative of $\mathbf{u}(\zeta \mathbf{a}'')$ in ζ evaluated at $\zeta = 0$. This derivative is strictly positive in every entry, since (by convexity of the u_i) the entry v_i must exceed $(u_i(\mathbf{a}'') - u_i(\mathbf{0}))/a''_i$. By the chain rule, $\mathbf{v} = \mathbf{J}(\mathbf{0})\mathbf{a}''$. From the fact that \mathbf{v} is positive, we deduce via simple algebraic manipulation that there is a positive vector \mathbf{w} so that $\mathbf{B}(\mathbf{0})\mathbf{w} > \mathbf{w}$. And from this it follows by the Collatz-Wielandt formula (Meyer 2000, eq. 8.3.3) that the spectral radius of $\mathbf{B}(\mathbf{0})$ exceeds 1.

Now consider the case in which assumption 3 does not hold.

First, suppose that $r(\mathbf{B}(\mathbf{0})) > 1$. Then the same is true when $\mathbf{B}(\mathbf{0})$ is replaced by one of its irreducible blocks, and in that case a Pareto improvement on $\mathbf{0}$ (involving only the agents in the irreducible block taking positive effort) is found as above. So $\mathbf{0}$ is not Pareto efficient.

⁴⁵ See statement 3 of the theorem in app. A.

⁴⁶ Suppose otherwise, and let \mathbf{a}'' be chosen so that $\mathbf{u}(\mathbf{a}'') - \mathbf{u}(\mathbf{0}) \geq \mathbf{0}$ (note that this is possible, since $\mathbf{a}'' = \mathbf{a}'$ satisfies this inequality) and so that the number of 0 entries in $\mathbf{u}(\mathbf{a}'') - \mathbf{u}(\mathbf{0})$ is as small as possible. Let S be the set of i for which $u_i(\mathbf{a}'') - u_i(\mathbf{a}) > 0$. Then by irreducibility of benefits, we can find $j \in S$ and $k \notin S$ such that $J_{kj}(\mathbf{0}) > 0$. Define $a'''_j = a''_j + \varepsilon$ and $a'''_i = a''_i$ for all $i \neq j$. If $\varepsilon > 0$ is chosen small enough, then by continuity of the u_i we have $u_i(\mathbf{a}''') - u_i(\mathbf{a}) > 0$ for all $i \in S$, but also $u_k(\mathbf{a}''') - u_k(\mathbf{a}) > 0$, contradicting the choice of \mathbf{a}' .

Conversely, suppose that $\mathbf{0}$ is not Pareto efficient. There is an $\mathbf{a}' \in \mathbb{R}_+^n$ such that $u_i(\mathbf{a}') \geq u_i(\mathbf{0})$ for each i , with strict inequality for some i . Let $P = \{i: a'_i > 0\}$ be the set of agents taking positive actions at \mathbf{a}' . And let $\hat{\mathbf{B}}(\mathbf{0})$ be obtained by restricting $\mathbf{B}(\mathbf{0})$ to P (i.e., by throwing away rows and columns not corresponding to indices in P). For each $i \in P$, there is a $j \in P$ such that $\hat{B}_{ij}(\mathbf{0}) > 0$; otherwise, i would be worse off than at $\mathbf{0}$. Therefore, there is a cycle in $\hat{\mathbf{B}}(\mathbf{0})$.⁴⁷ And it follows, shrinking P if necessary, that $\hat{\mathbf{B}}(\mathbf{0})$ is irreducible when restricted to P . Next, applying the argument of footnote 46, we can find \mathbf{a}'' such that $u_i(\mathbf{a}'') > u_i(\mathbf{0})$ for each $i \in P$. From this point, we can argue as above⁴⁸ to conclude that $r(\hat{\mathbf{B}}(\mathbf{0})) > 1$. Since $\hat{\mathbf{B}}(\mathbf{0})$ is a submatrix of $\mathbf{B}(\mathbf{0})$, by fact 1, $r(\mathbf{B}(\mathbf{0})) > 1$. QED

Proof of Theorem 1

We first prove the following lemma.

LEMMA 2. If $\mathbf{a}^* \neq \mathbf{0}$ is a Lindahl outcome for preference profile \mathbf{u} , then $\mathbf{a}^* \in \mathbb{R}_{++}^n$.

Proof. Assume, toward a contradiction, that \mathbf{a}^* has some entries equal to 0. Let \mathbf{P} be a matrix of prices that support \mathbf{a}^* as a Lindahl outcome. Let S be the set of i so that $a_i^* = 0$, which is a proper subset of N , since $\mathbf{a}^* \neq \mathbf{0}$. By assumption 3 (connectedness of benefit flows), there is an $i \in S$ and a $j \notin S$ so that $J_{ij}(\mathbf{a}^*) > 0$. We argue that this implies

$$P_{ij} > 0.$$

If this were not true, then an $\mathbf{a} \neq \mathbf{a}^*$ in which only j increases his action slightly relative to \mathbf{a}^* would satisfy inequality $(\widehat{\mathbf{BB}}_i(\mathbf{P}))$ in definition 1 and be preferred by i to the outcome \mathbf{a}^* , contradicting the definition of a Lindahl outcome.

Now consider $(\widehat{\mathbf{BB}}_i(\mathbf{P})')$, the budget balance condition of agent i , at the outcome \mathbf{a}^* :

$$\sum_{k: k \neq i} P_{ik} a_k^* \leq a_i^* \sum_{k: k \neq i} P_{ki}.$$

Since $a_i^* = 0$, the right-hand side of this is 0. But $P_{ij} > 0$, and $a_j^* > 0$ (since $j \notin S$), so the left-hand side is positive. That is a contradiction. QED

It is now convenient to use an equivalent definition of Lindahl outcomes.

DEFINITION 5. An action profile \mathbf{a}^* is a Lindahl outcome for a preference profile \mathbf{u} if there exists an n -by- n matrix \mathbf{P} with each column summing to 0, so that the following conditions hold for every i : (a) the inequality

$$\sum_{j \in N} P_{ij} a_j \leq 0 \quad (\widehat{\mathbf{BB}}_i(\mathbf{P}))$$

is satisfied when $\mathbf{a} = \mathbf{a}^*$; and (b) for any \mathbf{a} such that inequality $(\widehat{\mathbf{BB}}_i(\mathbf{P}))$ is satisfied, we have $u_i(\mathbf{a}^*) \geq u_i(\mathbf{a})$.

⁴⁷ Recall the definition in Sec. III.B.

⁴⁸ The Collatz-Wielandt formula does not assume irreducibility.

Given a Lindahl outcome defined as in definition 1, set $P_{ii} = -\sum_{j:j \neq i} P_{ji}$ to find prices satisfying the new definition.⁴⁹ Conversely, the prices of definition 5 work in definition 1 without modification, since the original definition does not involve the diagonal terms of \mathbf{P} at all.

We now show that b implies a. Suppose that $\mathbf{a}^* \in \mathbb{R}_+^n$ is a nonzero Lindahl outcome. Lemma 2 implies that $\mathbf{a}^* \in \mathbb{R}_{++}^n$ or, in other words, that \mathbf{a}^* has only positive entries. Let \mathbf{P} be the matrix of prices satisfying the conditions of definition 5. Consider the following program for each $i \in N$, denoted by $\Pi_i(\mathbf{P})$:

$$\text{maximize } u_i(\mathbf{a}) \text{ subject to } \mathbf{a} \in \mathbb{R}_+^n \text{ and } \widehat{\mathbf{B}\mathbf{B}}_i(\mathbf{P}).$$

By definition of a Lindahl outcome, \mathbf{a}^* solves $\Pi_i(\mathbf{P})$. By assumption 3, there is some agent $j \neq i$ such that increases in his action a_j would make i better off. Arguing as in the proof of lemma 2, it holds that $P_{ij} > 0$. Therefore, the budget balance constraint $\widehat{\mathbf{B}\mathbf{B}}_i(\mathbf{P})$ is satisfied with equality, so that $\mathbf{P}\mathbf{a}^* = \mathbf{0}$. Because \mathbf{a}^* is interior, the gradient of the maximand u_i must be orthogonal to the constraint set given by inequality $\widehat{\mathbf{B}\mathbf{B}}_i(\mathbf{P})$. In other words, row i of $\mathbf{J}(\mathbf{a}^*)$ is parallel to row i of \mathbf{P} . These facts together imply that $\mathbf{J}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{0}$ and so $\mathbf{B}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{a}^*$ (see Sec. IV.B.1).

We now show that a implies b. Since \mathbf{a}^* is a nonnegative right-hand eigenvector of $\mathbf{B}(\mathbf{a}^*)$, the Perron-Frobenius theorem guarantees that 1 is a largest eigenvalue of $\mathbf{B}(\mathbf{a}^*)$. Arguing as in the proof of proposition 1a, we deduce that there is a nonzero vector $\boldsymbol{\theta}$ for which $\boldsymbol{\theta}\mathbf{J}(\mathbf{a}^*) = \mathbf{0}$. We need to find prices supporting \mathbf{a}^* as a Lindahl outcome. Define the matrix \mathbf{P} by $P_{ij} = \theta_j J_{ij}(\mathbf{a}^*)$, and note that for all $j \in N$ we have

$$\sum_{i \in N} P_{ij} = \sum_{i \in N} \theta_j J_{ij}(\mathbf{a}^*) = (\boldsymbol{\theta}\mathbf{J}(\mathbf{a}^*))_j = 0, \quad (\text{C1})$$

where $(\boldsymbol{\theta}\mathbf{J}(\mathbf{a}^*))_j$ refers to entry j of the vector $\boldsymbol{\theta}\mathbf{J}(\mathbf{a}^*)$.

Note that $\mathbf{B}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{a}^*$ implies $\mathbf{J}(\mathbf{a}^*)\mathbf{a}^* = \mathbf{0}$ and that each row of \mathbf{P} is just a scaling of the corresponding row of $\mathbf{J}(\mathbf{a}^*)$. We therefore have

$$\mathbf{P}\mathbf{a}^* = \mathbf{0}. \quad (\text{C2})$$

We claim that, for each i , the vector \mathbf{a}^* solves $\Pi_i(\mathbf{P})$. This is because the gradient of u_i at \mathbf{a}^* , which is row i of $\mathbf{J}(\mathbf{a}^*)$, is normal to the constraint set by construction of \mathbf{P} . Moreover, by equation (C2), \mathbf{a}^* satisfies the constraint $(\widehat{\mathbf{B}\mathbf{B}}_i(\mathbf{P}))$. The claim then follows by the concavity of u_i . QED

Proof of Proposition 2

We use the Kakutani fixed-point theorem to find a centrality action profile. Define $Y = \{\mathbf{a} \in \mathbb{R}_+^n : \min_i [\mathbf{J}(\mathbf{a})\mathbf{a}]_i > 0\}$, the set of action profiles \mathbf{a} at which every-

⁴⁹ In essence, $-P_{ii}$ is the total subsidy agent i receives per unit of effort, equal to the sum of personalized taxes paid by other people to him for his effort.

one has positive gains from scaling \mathbf{a} up. It is easy to check that Y is convex.⁵⁰ Also, Y is bounded by assumption 4. Thus, \bar{Y} , the closure of Y , is compact.⁵¹

Define the correspondence $F: \bar{Y} \setminus \{0\} \rightrightarrows \bar{Y}$ by

$$F(\mathbf{a}) = \left\{ \lambda \mathbf{a} \in \bar{Y} : \lambda \geq 0 \text{ and } \min_i (\mathbf{J}(\lambda \mathbf{a}) \mathbf{a})_i \leq 0 \right\}.$$

This correspondence, at an argument \mathbf{a} , has in its image all actions $\lambda \mathbf{a}$ (i.e., on the same ray as \mathbf{a}) such that, at $\lambda \mathbf{a}$, at least one agent does not want to further scale up actions. Finally, recalling the definition of $\mathbf{d}^{\text{eg}}(\mathbf{a})$ from appendix B, define the correspondence $G: \bar{Y} \rightrightarrows \bar{Y}$ by

$$G(\mathbf{a}) = F(\mathbf{d}^{\text{eg}}(\mathbf{a})).$$

Note that $\mathbf{d}^{\text{eg}}(\mathbf{a})$ is always nonzero, so that the argument of F is in its domain.⁵² The function \mathbf{d}^{eg} is continuous (Wilkinson 1965, 66–67), and F is clearly upper hemicontinuous, so it follows that G is upper hemicontinuous. Finally, from the definitions of Y and F it follows that F is nonempty-valued.⁵³ Since \bar{Y} is a compact and convex set, the Kakutani fixed-point theorem implies that there is an $\mathbf{a} \in \bar{Y}$ such that $\mathbf{a} \in F(\mathbf{d}^{\text{eg}}(\mathbf{a}))$. Writing $\hat{\mathbf{a}} = \mathbf{d}^{\text{eg}}(\mathbf{a})$, this means that there is some $\lambda \geq 0$ such that $\min_i (\mathbf{J}(\lambda \hat{\mathbf{a}}) \hat{\mathbf{a}})_i \leq 0$. Let $\mathbf{a}^* = \lambda \hat{\mathbf{a}}$. We argue that \mathbf{a}^* satisfies scaling indifference (and is therefore a centrality action profile).

Suppose that $\mathbf{a}^* \neq 0$. Then by continuity of the function $\lambda \mapsto \mathbf{J}(\lambda \hat{\mathbf{a}}) \hat{\mathbf{a}}$, there is some i for which we have $[\mathbf{J}(\mathbf{a}^*) \hat{\mathbf{a}}]_i = 0$, so that some player's marginal benefit to scaling is equal to his marginal cost. Since $\hat{\mathbf{a}}$ is an egalitarian direction at the action profile \mathbf{a}^* , the equation $[\mathbf{J}(\mathbf{a}^*) \hat{\mathbf{a}}]_i = 0$ must hold for all i , and therefore $\mathbf{J}(\mathbf{a}^*) \hat{\mathbf{a}} = 0$. Since $\hat{\mathbf{a}}$ and \mathbf{a}^* are scalings of each other, we deduce that $\mathbf{J}(\mathbf{a}^*) \mathbf{a}^* = 0$. The condition $\mathbf{J}(\mathbf{a}^*) \mathbf{a}^* = 0$ and assumption 3—connectedness of benefit flows—imply that $\mathbf{a}^* \in \mathbb{R}_{++}^n$.

If $\mathbf{a}^* = 0$, consider the bang-for-the-buck vector $\mathbf{b}(0, \hat{\mathbf{a}})$, which corresponds to starting at 0 and moving in the egalitarian direction $\hat{\mathbf{a}}$. Because $\hat{\mathbf{a}}$ is egalitarian, we can write $\mathbf{b}(0, \hat{\mathbf{a}}) = b\mathbf{1}$ for some b . And we can deduce that b is no greater than

⁵⁰ Given $\mathbf{a}, \mathbf{a}' \in Y$, and $\lambda \in [0, 1]$, define $\mathbf{a}'' = \lambda \mathbf{a} + (1 - \lambda) \mathbf{a}'$. Note that for all $i \in N$ and $\varepsilon \geq -1$,

$$u_i((1 + \varepsilon) \mathbf{a}'') \geq \lambda u_i((1 + \varepsilon) \mathbf{a}) + (1 - \lambda) u_i((1 + \varepsilon) \mathbf{a}')$$

by concavity of the u_i . Differentiating in ε at $\varepsilon = 0$ yields the result.

⁵¹ It is tempting to define $Y = \{ \mathbf{a} \in \mathbb{R}_+^n : \min_i (\mathbf{J}(\mathbf{a}) \mathbf{a})_i \geq 0 \}$ instead and avoid having to take closures, but this set can be unbounded even when \bar{Y} as we defined it above is bounded. For example, our assumptions do not exclude the existence of an (infinite) ray along which $\min_i (\mathbf{J}(\mathbf{a}) \mathbf{a})_i = 0$.

⁵² Even though the domain of F is not a compact set, G is a correspondence from a compact set into itself.

⁵³ Toward a contradiction, take a nonzero \mathbf{a} such that $F(\mathbf{a})$ is empty. Let $\bar{\lambda}$ be the maximum λ such that $\lambda \mathbf{a} \in \bar{Y}$; such a $\bar{\lambda}$ exists because \mathbf{a} is nonzero and \bar{Y} is compact. Since $\mathbf{J}(\bar{\lambda} \mathbf{a}) \mathbf{a} > 0$, it follows that for all i ,

$$\left. \frac{du_i((1 + \varepsilon) \bar{\lambda} \mathbf{a})}{d\varepsilon} \right|_{\varepsilon=0} > 0,$$

from which it follows that $(\bar{\lambda} + \delta) \mathbf{a} \in Y$ for small enough δ . This contradicts the choice of $\bar{\lambda}$ (recalling the definition of Y).

1—otherwise, $F(\hat{\mathbf{a}})$ would not contain $\mathbf{a}^* = \mathbf{0}$. By proposition 6, it follows that $r(\mathbf{B}(\mathbf{0})) \leq 1$. Then, $\mathbf{0}$ is Pareto efficient by proposition 1b. QED

Proof of Proposition 3

Let $\mathcal{W}_i^\ell(\ell; \mathbf{M})$ be the set of all walks of length ℓ in a matrix \mathbf{M} starting at i , so that $w(1) = i$. The proof follows immediately from the following observation.

FACT 2. For any irreducible, nonnegative, aperiodic matrix \mathbf{Q} and any i, j ,

$$\frac{p_i}{p_j} = \lim_{\ell \rightarrow \infty} \frac{\sum_{w \in \mathcal{W}_i^\ell(\ell; \mathbf{Q})} v(w; \mathbf{Q})}{\sum_{w \in \mathcal{W}_j^\ell(\ell; \mathbf{Q})} v(w; \mathbf{Q})},$$

where \mathbf{p} is any nonnegative right-hand eigenvector of \mathbf{Q} (i.e., a right-hand Perron vector in the terminology of app. A).

Proof. Note that the formula above is equivalent to

$$\frac{p_i}{p_j} = \lim_{\ell \rightarrow \infty} \frac{\sum_k [\mathbf{Q}^\ell]_{ik}}{\sum_k [\mathbf{Q}^\ell]_{jk}}, \quad (\text{C3})$$

where $[\mathbf{Q}^\ell]_{ik}$ denotes the entry in the (i, k) position of the matrix \mathbf{Q}^ℓ . To prove equation (C3), let $\rho = r(\mathbf{Q})$ and note that

$$\lim_{\ell \rightarrow \infty} \left(\frac{\mathbf{Q}}{\rho} \right)^\ell = \mathbf{w}^\top \mathbf{p}, \quad (\text{C4})$$

where \mathbf{w} is a left-hand Perron vector of \mathbf{Q} and \mathbf{p} is a right-hand Perron vector (recall app. A). This is statement (8.3.13) in Meyer (2000); the hypothesis that \mathbf{Q} is primitive in that statement follows from the assumed aperiodicity of \mathbf{Q} (see theorems 1 and 2 of Perkins [1961]). To conclude, observe that equation (C4) directly implies equation (C3). QED

To prove the proposition from fact 2, set $\mathbf{Q} = \mathbf{B}(\mathbf{a}) = \mathbf{M}^\top$ and note that then the right-hand side of the equation in fact 2 is equal to the right-hand side of the equation in proposition 3. The statement that \mathbf{a} has the centrality property is equivalent to the statement that \mathbf{a} is a right-hand Perron eigenvector of $\mathbf{Q} = \mathbf{B}(\mathbf{a})$.

Proof of Lemma 1

Suppose that this does not hold, and let M be the nonempty set of all i such that $J_{ii}(\mathbf{a}^*)\theta_i = 0$. By assumption 3, if M is not the set of all agents, there is some $i \in M$ and $j \notin M$ with $J_{ij}(\mathbf{a}^*) > 0$, which implies $\theta_j = 0$, a contradiction to the definition of M . If M is the set of all agents, then let $\tilde{\mathbf{J}}$ be equal to $\mathbf{J}(\mathbf{a}^*; \hat{\mathbf{u}})$ with the diagonal zeroed out, and note that $\tilde{\mathbf{J}}$ is an irreducible, nonnegative matrix with $\tilde{\theta}\tilde{\mathbf{J}} = \mathbf{0}$, again a contradiction (since θ was assumed to be nonzero).

Proof of Proposition 5

For $j \in M$, set

$$m_j(\mathbf{a}) = \theta_j^{-1} \sum_{i \notin M} \theta_i J_{ij}(\mathbf{a}^*) a_j.$$

One can check that with these payments, the problem of maximizing

$$\sum_{j \in M} \theta_j \tilde{u}_j(\mathbf{a})$$

has the same first-order conditions evaluated at \mathbf{a}^* as the planner's problem in the grand coalition; these are

$$\sum_{i \in N} \theta_i J_{ij}(\mathbf{a}^*) = 0$$

for each j . So the social planner's problem in group M , of maximizing the weighted sum of utilities is solved by $\mathbf{a} = \mathbf{a}^*$. Because the utility functions are concave, the solution is, indeed, Pareto efficient for \mathbf{a}^* . The analogous argument holds for M^c .

Appendix D

Additional Results

PROPOSITION 7. The following are equivalent: (a) $r(\mathbf{B}(\mathbf{0})) \leq 1$; (b) $\mathbf{0}$ is a Pareto efficient action profile; and (c) $\mathbf{0}$ is a Lindahl outcome.

Proof. Proposition 1b establishes the equivalence between a and b.

$b \Rightarrow c$.—The construction of prices is exactly analogous to the proof of theorem 1; the only difference is that, rather than the Pareto weights, we use Pareto weights adjusted by the Lagrange multipliers on the binding constraints $a_i \geq 0$.

$c \Rightarrow b$.—The standard proof of the first welfare theorem goes through without modification; see, for example, Foley (1970). QED

Appendix E

Essential versus Key Players

In this appendix, we compare the concept of a key player from Ballester et al. (2006) with our concept of essential agents, as defined in Section III.B. Suppose that there are four agents with the following utility functions:

$$u_1 = 10a_3 - 0.5a_1^2 + a_1a_4,$$

$$u_2 = 10a_1 - 0.5a_2^2 + a_2a_4,$$

$$u_3 = 10a_2 - 0.5a_3^2 + a_3a_4,$$

$$u_4 = a_4 - 0.5a_4^2.$$

If all agents take actions greater than zero, these utilities induce the benefits network shown in figure E1, where an arrow from i to j means $\partial u_i / \partial a_j > 0$.

It is easy to see that the unique Nash equilibrium is $a_1^* = a_2^* = a_3^* = a_4^* = 1$. Following the exercise in Ballester et al. (2006), suppose that player 4 is "removed" from the network, so that player 4 becomes disconnected, providing no benefits to the other players. Formally, in all utility functions other than that of player 4, the action a_4 is set to zero:

$$u_1 = 10a_3 - 0.5a_1^2,$$

$$u_2 = 10a_1 - 0.5a_2^2,$$

$$u_3 = 10a_2 - 0.5a_3^2,$$

$$u_4 = a_4 - 0.5a_4^2.$$

In this new network, the unique Nash equilibrium is $a_1^* = a_2^* = a_3^* = 0$, while $a_4^* = 1$. Therefore, the removal of player 4 decreases the actions by all other players. Suppose instead that player $i \neq 4$ were removed. In the network the unique Nash equilibrium is $a_j^* = 1$ for all $j \neq i$ and $x_i = 0$. Thus, the action profile after player 4 is removed is pointwise dominated by the action profile after any other player is removed. And so aggregate actions decrease the most when player 4 is removed, which implies that player 4 the key player, as defined by Ballester et al. (2006).

Consider now whether a Pareto improvement is possible at the Nash equilibrium action profile. By proposition 1 this is possible if and only if the spectral radius of the benefits matrix is greater than 1. In the essential-player exercise, we remove a player from negotiations by having that player take his status quo action—in this case, his Nash equilibrium action. Because this player is unable to provide positive marginal benefits to anyone else, we remove him from the benefits network when looking for Pareto improvements. In the subnetwork without player 4, the spectral radius is greater than 1, and so a Pareto improvement is possible and player 4 is not an essential player. However, in the subnetwork without any player $i \neq 4$, the spectral radius is 0, and so a Pareto improvement is not possible. Thus, all players $i \neq 4$ are essential.

In summary, player 4 is the key player, while all other players, and not player 4, are essential. What makes a player “key” in Ballester et al. (2006) is the complementarity of his action with the actions of others. Player 4 is the only player with such complementarities, since the other players $i \neq 4$ have terms $x_i x_4$ in their util-




FIG. E1.—The benefits matrix in our example

ity functions. In contrast, what makes a player “essential” is his position in cycles in the benefits network. When any $i \neq 4$ is removed, the benefits network has no cycles, and so such players are essential. In contrast, when player 4 is removed, there is still a strong cycle among the remaining players in the benefits network.

References

- Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. “The Network Origins of Aggregate Fluctuations.” *Econometrica* 80 (5): 1977–2016.
- Acemoglu, Daron, Camilo García-Jimeno, and James A. Robinson. 2015. “State Capacity and Economic Development: A Network Approach.” *A.E.R.* 105:2364–2409.
- Allouch, Nizar. 2015. “On the Private Provision of Public Goods on Networks.” *J. Econ. Theory* 157:527–52.
- . 2017. “The Cost of Segregation in (Social) Networks.” *Games and Econ. Behavior* 106:329–42. <http://www.sciencedirect.com/science/article/pii/S089982561730146X>.
- Altman, Alon, and Moshe Tennenholtz. 2005. “Ranking Systems: The PageRank Axioms.” In *EC '05: Proceedings of the 6th ACM Conference on Electronic Commerce*, 1–8. New York: Association for Computing Machinery. <http://doi.acm.org/10.1145/1064009.1064010>.
- Arrow, Kenneth J. 1969. “The Organization of Economic Activity: Issues Pertinent to the Choice of Market versus Non-market Allocation.” In *The Analysis and Evaluation of Public Expenditures: The PBB-System: A Compendium of Papers Submitted to the Subcommittee on Economy in Government of the Joint Economic Committee*, vol. 1. Washington, DC: Government Printing Office.
- Baetz, Oliver. 2015. “Social Activity and Network Formation.” *Theoretical Econ.* 10 (2): 315–40.
- Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou. 2006. “Who’s Who in Networks. Wanted: The Key Player.” *Econometrica* 74 (5): 1403–17.
- Bergstrom, Theodore, Lawrence Blume, and Hal Varian. 1986. “On the Private Provision of Public Goods.” *J. Public Econ.* 29 (1): 25–49.
- Bonacich, Phillip B. 1987. “Power and Centrality: A Family of Measures.” *American J. Sociology* 92:1170–82.
- Bramoullé, Yann, and Rachel Kranton. 2007. “Public Goods in Networks.” *J. Econ. Theory* 135 (1): 478–94.
- Bramoullé, Yann, Rachel Kranton, and Martin D’Amours. 2014. “Strategic Interaction and Networks.” *A.E.R.* 104 (3): 898–930.
- Brin, Sergey, and Lawrence Page. 1998. “The Anatomy of a Large-Scale Hypertextual Web Search Engine.” *Computer Networks and ISDN Systems* 30:107–17. <http://dl.acm.org/citation.cfm?id=297805.297827>.
- Calvó-Armengol, Antoni, Eleonora Patacchini, and Yves Zenou. 2009. “Peer Effects and Social Networks in Education.” *Rev. Econ. Studies* 76 (4): 1239–67.
- Chen, Ying-Ju, Yves Zenou, and Junjie Zhou. 2018. “Competitive Pricing Strategies in Social Networks.” *RAND J. Econ.* 49 (3): 672–705.
- Conitzer, Vincent, and Tuomas Sandholm. 2004. “Expressive Negotiation over Donations to Charities.” In *EC '04: Proceedings of the 5th ACM conference on Electronic Commerce*, 51–60. New York: Assoc. Computing Machinery.
- Dávila, Julio, and Jan Eeckhout. 2008. “Competitive Bargaining Equilibrium.” *J. Econ. Theory* 139 (1): 269–94.

- Dávila, Julio, Jan Eeckhout, and César Martinelli. 2009. "Bargaining over Public Goods." *J. Public Econ. Theory* 11 (6): 927–45.
- Dequiedt, Vianney, and Yves Zenou. 2017. "Local and Consistent Centrality Measures in Parameterized Networks." *Math. Soc. Sci.* 88:28–36.
- Du, Ye, Ehud Lehrer, and Ady Pauzner. 2015. "Competitive Economy as a Ranking Device over Networks." *Games and Econ. Behavior* 91:1–13.
- Elliott, Matthew, and Benjamin Golub. 2015. "Ranking Agendas for Negotiations." Manuscript, Dept. Econ., Harvard Univ.
- Foley, Duncan K. 1970. "Lindahl's Solution and the Core of an Economy with Public Goods." *Econometrica* 38 (1): 66–72.
- Friedkin, Noah E. 1991. "Theoretical Foundations for Centrality Measures." *American J. Sociology* 96:1478–1504.
- Galeotti, Andrea, and Sanjeev Goyal. 2010. "The Law of the Few." *A.E.R.* 100:1468–92.
- Garratt, Rod, and Marek Pycia. 2016. "Efficient Bilateral Trade." Manuscript, Dept. Econ., Univ. California, Los Angeles.
- Ghosh, Arpita, and Mohammad Mahdian. 2008. "Charity Auctions on Social Networks." In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1019–28. Philadelphia: Soc. Indus. and Appl. Math.
- Hartfiel, D. J., and Carl D. Meyer. 1998. "On the Structure of Stochastic Matrices with a Subdominant Eigenvalue near 1." *Linear Algebra and Its Applications* 272:193–203.
- Hatfield, John William, Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp. 2013. "Stability and Competitive Equilibrium in Trading Networks." *J.P.E.* 121 (5): 966–1005.
- Hiller, Timo. 2017. "Peer Effects in Endogenous Networks." *Games and Econ. Behavior* 105:349–67.
- Hurwicz, Leonid. 1979a. "On Allocations Attainable through Nash Equilibria." *J. Econ. Theory* 21 (1): 140–65.
- . 1979b. "Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points." *Rev. Econ. Studies* 46:217–25.
- Jackson, Matthew O. 2008. *Social and Economic Networks*. Princeton, NJ: Princeton Univ. Press.
- Johnson, Charles R., and Pablo Tarazaga. 2004. "On Matrices with Perron-Frobenius Properties and Some Negative Entries." *Positivity* 8 (4): 327–38. doi:10.1007/s11117-003-3881-3.
- Katz, Leo. 1953. "A New Status Index Derived from Sociometric Analysis." *Psychometrika* 18:39–43.
- Kendall, Maurice George. 1955. "Further Contributions to the Theory of Paired Comparisons." *Biometrics* 11 (1): 43–62.
- Langville, Amy N., and Carl D. Meyer. 2012. *Google's PageRank and Beyond: The Science of Search Engine Rankings*. Princeton, NJ: Princeton Univ. Press.
- McKenzie, Lionel. 1960. "Matrices with Dominant Diagonals and Economic Theory." In *Mathematical Methods in the Social Sciences, 1959: Proceedings of the First Stanford Symposium*, edited by Kenneth Joseph Arrow, Samuel Karlin, and Patrick Suppes, 47–62. Stanford, CA: Stanford Univ. Press.
- Meyer, Carl D. 2000. *Matrix Analysis and Applied Linear Algebra*. Philadelphia: Soc. Indus. and Appl. Math.
- Milnor, John W. 2001. "Appendix: Matrix Estimates and the Perron-Frobenius Theorem." Manuscript, Dept. Mathematics, Stony Brook Univ. <http://www.math.stonybrook.edu/~jack/DYNOTES/dnA.pdf>.

- Noutsos, Dimitrios. 2006. "On Perron-Frobenius Property of Matrices Having Some Negative Entries." *Linear Algebra and Its Applications* 412 (2–3): 132–53.
- Palacios-Huerta, Ignacio, and Oscar Volij. 2004. "The Measurement of Intellectual Influence." *Econometrica* 72 (3): 963–77.
- Penta, Antonio. 2011. "Multilateral Bargaining and Walrasian Equilibrium." *J. Math. Econ.* 47:417–24.
- Perkins, Peter. 1961. "A Theorem on Regular Matrices." *Pacific J. Math.* 11 (4): 1529–33.
- Prat, Andrea, Joan de Martí, and Antoni Calvó-Armengol. 2015. "Communication and Influence." *Theoretical Econ.* 10:649–90.
- Pycia, Marek, and M. Bumin Yenmez. 2015. "Matching with Externalities." Manuscript, Dept. Econ., Univ. California, Los Angeles. <https://ssrn.com/abstract=2475468>.
- Samuelson, Paul A. 1954. "The Pure Theory of Public Expenditure." *Rev. Econ. and Statis.* 36 (4): 387–89.
- Shannon, Chris. 2008. "Determinacy and Indeterminacy of Equilibria." In *The New Palgrave Dictionary of Economics*, edited by Steven N. Durlauf and Lawrence E. Blume, 451–56. Basingstoke: Palgrave Macmillan.
- Shapley, Lloyd S., and Martin Shubik. 1969. "On the Core of an Economic System with Externalities." *A.E.R.* 59 (4): 678–84.
- Takayama, Akira. 1985. *Mathematical Economics*. 2nd ed. Cambridge: Cambridge Univ. Press.
- von Luxburg, Ulrike. 2007. "A Tutorial on Spectral Clustering." *Statis. and Computing* 17 (4): 395–416.
- Wei, Teh-Hsing. 1952. *The Algebraic Foundations of Ranking Theory*. London: Cambridge Univ. Press.
- Wilkinson, J. H. 1965. *The Algebraic Eigenvalue Problem*. Oxford: Oxford Univ. Press.
- Yildiz, Muhamet. 2003. "Walrasian Bargaining." *Games and Econ. Behavior* 45 (2): 465–87.