

How Bad is Forming Your Own Opinion?

David Bindel
Dept. of Computer Science
Cornell University
Ithaca NY 14853
Email: bindel@cs.cornell.edu

Jon Kleinberg
Dept. of Computer Science
Cornell University
Ithaca NY 14853
Email: kleinber@cs.cornell.edu

Sigal Oren
Dept. of Computer Science
Cornell University
Ithaca NY 14853
Email: sigal@cs.cornell.edu

Abstract— A long-standing line of work in economic theory has studied models by which a group of people in a social network, each holding a numerical opinion, can arrive at a shared opinion through repeated averaging with their neighbors in the network. Motivated by the observation that consensus is rarely reached in real opinion dynamics, we study a related sociological model in which individuals' intrinsic beliefs counterbalance the averaging process and yield a diversity of opinions.

By interpreting the repeated averaging as best-response dynamics in an underlying game with natural payoffs, and the limit of the process as an equilibrium, we are able to study the cost of disagreement in these models relative to a social optimum. We provide a tight bound on the cost at equilibrium relative to the optimum; our analysis draws a connection between these agreement models and extremal problems for generalized eigenvalues. We also consider a natural network design problem in this setting, where adding links to the underlying network can reduce the cost of disagreement at equilibrium.

1. INTRODUCTION

Averaging Opinions in a Social Network: An active line of recent work in economic theory has considered processes by which a group of people connected in a social network can arrive at a shared opinion through a form of repeated averaging [8], [11], [13]. This work builds on a basic model of DeGroot [7], in which we imagine that each person i holds an *opinion* equal to a real number z_i , which might for example represent a position on a political spectrum, or a probability that i assigns to a certain belief. There is a weighted graph $G = (V, E)$ representing a social network, and node i is influenced by the opinions of her neighbors in G , with the edge weights reflecting the extent of this influence. Thus, in each time step node i updates her opinion to be a weighted average of her current opinion with the current opinions of her neighbors.

This body of work has developed a set of general conditions under which such processes will converge to

a state of *consensus*, in which all nodes hold the same opinion. This emphasis on consensus, however, restricts the focus of the modeling activity to a specific type of opinion dynamics, where the opinions of the group all come together. As the sociologist David Krackhardt has observed,

We should not ignore the fact that in the real world consensus is usually not reached. Recognizing this, most traditional social network scientists do not focus on an equilibrium of consensus. They are instead more likely to be concerned with explaining the lack of consensus (the variance) in beliefs and attitudes that appears in actual social influence contexts [15].

In this paper we study a model of opinion dynamics in which consensus is not in general reached, and our main goal is to quantify the inherent social cost of this lack of consensus. To do this, we first need a framework that captures some of the underlying reasons why consensus is not reached, as well as a way of measuring the cost of disagreement.

Lack of Agreement and its Cost: We begin from a variation on the DeGroot model due to Friedkin and Johnsen [10], which posits that each node i maintains a persistent *internal opinion* s_i , which remains constant even as node i updates her overall opinion z_i through averaging. More precisely, if $w_{i,j} \geq 0$ denotes the weight on the edge (i, j) in G , then in one time step node i updates her opinion to be the average

$$z_i = \frac{s_i + \sum_{j \in N(i)} w_{i,j} z_j}{1 + \sum_{j \in N(i)} w_{i,j}}, \quad (1)$$

where $N(i)$ denotes the set of neighbors of i in G . Note that because of the presence of s_i as a constant in each iteration, repeated averaging will not in general bring all nodes to the same opinion. In this way, the model distinguishes between an individual's intrinsic belief s_i and her overall opinion z_i ; the latter represents a compromise between the persistent value of s_i and the expressed opinions of others to whom i is connected. This distinction between s_i and z_i also has parallels in empirical work that seeks to trace deeply held opinions

A full version of this paper is available at the authors' Web pages. The research has been supported in part by the MacArthur Foundation, the Sloan Foundation, a Google Research Grant, a Yahoo! Research Alliance Grant, and NSF grants IIS-0910664, CCF-0910940, and IIS-1016099.

such as political orientations back to differences in education and background, and even to explore genetic bases for such patterns of variation [2].

Now, if consensus is not reached, how should we quantify the cost of this lack of consensus? Here we observe that since the standard models use averaging as their basic mechanism, we can equivalently view nodes' actions in each time step as myopically optimizing a quadratic cost function: updating z_i as in Equation (1) is the same as choosing z_i to minimize

$$(z_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j} (z_i - z_j)^2. \quad (2)$$

We therefore take this as the *cost* that i incurs by choosing a given value of z_i , so that averaging becomes a form of cost minimization. Indeed, more strongly, we can think of repeated averaging as the trajectory of best-response dynamics in a game played by the nodes in V , where i 's strategy is a choice of opinion z_i , and her payoff is the negative of the cost in Equation (2).

Nash Equilibrium and Social Optimality in a Game of Opinion Formation: In this model, repeated averaging — while it does not in general converge to consensus among all nodes — does converge to the unique Nash equilibrium of the game defined by the individual cost functions in (2): each node i has an opinion x_i that is the weighted average of i 's internal opinion and the (equilibrium) opinions of i 's neighbors. This equilibrium will not in general correspond to the *social optimum*, the vector of node opinions $y = (y_i : i \in V)$ that minimizes the *social cost*, defined to be sum of all players' costs: $c(y) = \sum_i \left((y_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j} (y_i - y_j)^2 \right)$.

The sub-optimality of the Nash equilibrium can be viewed in terms of the *externality* created by a player's personal optimization: by refusing to move further toward their neighbors' opinions, players can cause additional cost to be incurred by these neighbors without accounting for it in their own objective function. In fact we can view the problem of minimizing social cost for this game as a type of *metric labeling problem* [5], [14], albeit a polynomial-time solvable one with a non-metric quadratic distance function on the real numbers: we seek node labels that balance the value of a cost function at each node (capturing disagreement with node-level information) and a cost function for label disagreement across edges. Viewed this way, the sub-optimality of Nash equilibrium becomes a kind of sub-optimality for local optimization.

A natural question is thus the *price of anarchy* for this basic model of opinion formation: how far from optimality can the Nash equilibrium be?

Our Results: Undirected Graphs: The model we have described can be used as stated in both undirected




Figure 1. An example in which the two players on the sides do not compromise by the optimal amount, given that the player in the middle should not shift her opinion. The social cost of the optimal set of opinions is $1/3$, while the cost of the Nash equilibrium is $3/8$.

and directed graphs — the only difference is in whether i 's neighbor set $N(i)$ represents the nodes to whom i is connected by undirected edges, or to whom i links with directed edges. However, the behavior of the price of anarchy is very different in undirected and directed graphs, and so we analyze them separately, beginning with the undirected case.

As an example of how a sub-optimal social cost can arise at equilibrium in an undirected graph, consider the graph depicted in Figure 1 — a three-node path in which the nodes have internal opinions 0, $1/2$, and 1 respectively. As shown in the figure, the ratio between the social cost of the Nash equilibrium and the social optimum is $9/8$. Intuitively, the reason for the higher cost of the Nash equilibrium is that the center node — by symmetry — cannot usefully shift her opinion in either direction, and so to achieve optimality the two outer nodes need to compromise more than they want to at equilibrium. This is a reflection of the externality discussed above, and it is the qualitative source of sub-optimality in general for equilibrium opinions — nodes move in the direction of their neighbors, but not sufficiently to achieve the globally minimum social cost.

Our first result is that the very simple example in Figure 1 is in fact extremal for undirected graphs: we show that for any undirected graph G and any vector of internal opinions s , the price of anarchy is at most $9/8$. We prove this by casting the question as an extremal problem for quadratic forms, and analyzing the resulting structure using eigenvalues of the Laplacian matrix of G . From this, we obtain a characterization of the set of graphs G for which some vector of internal opinions s yields a price of anarchy of $9/8$.

We show that this bound of $9/8$ continues to hold even for some generalizations of the model — when nodes i have different coefficients w_i on the cost terms for their internal opinions, and (in a kind of infinite limit of node weight) when certain nodes are “fixed” and simply do not modify their opinions.

Our Results: Directed Graphs: We next consider the case in which G is a directed graph; the form of the cost functions remains exactly the same, with

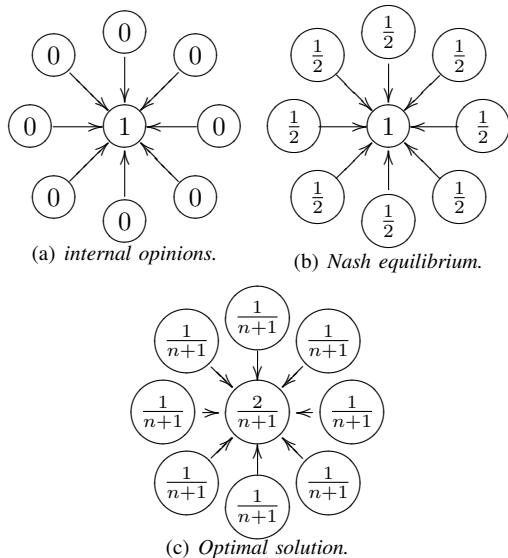


Figure 2. An example demonstrating the price of anarchy of a directed graph can be unbounded.

directed edges playing the role of undirected ones, but the range of possible behaviors in the model becomes very different, owing to the fact that nodes can now exert a large influence over the network without being influenced themselves. Indeed, as Matt Jackson has observed, directed versions of repeated averaging models can naturally incorporate “external” media sources; we simply include nodes with no outgoing links, so that in equilibrium they maintain precisely their internal opinion [13].

We first show that the spectral machinery developed for analyzing undirected graphs can be extended to the directed case; through an approach based on generalized eigenvalue problems we can efficiently compute the maximum possible price of anarchy, over all choices of internal node opinions, on a given graph G . However, unlike in the case of undirected graphs, this price of anarchy can be very large; the simple example in Figure 2 shows a case in which $n - 1$ nodes with internal opinion 0 all link to a single node that has internal opinion 1 and no out-going edges, producing an in-directed star. As a result, the social cost of the Nash equilibrium is $\Omega(n)$, whereas the minimum social cost is at most 1, since the player at the center of the star could simply shift her opinion to 0. Intuitively, this corresponds to a type of social network in which the whole group pays attention to a single influential “leader” or “celebrity”; this drags people’s opinions far from their internal opinions s_i , creating a large social cost. Unfortunately, the leader is essentially unaware of the people paying attention to her, and hence has no incentive to modify her opinion in a direction that could greatly reduce the social cost.

In Section 3 we show that a price of anarchy lower-bounded by a polynomial in n can in fact be achieved in directed graphs of constant degree, so this behavior is not simply a consequence of large in-degree. It thus becomes a basic question whether there are natural classes of directed graphs, and even bounded-degree directed graphs, for which a constant price of anarchy is achievable.

Unweighted Eulerian directed graphs are a natural class to consider — first, because they generalize undirected graphs, and second, because they capture the idea that at least at a local level no node has an asymmetric effect on the system. We use our analysis framework for directed graphs to derive bounds on the price of anarchy of two subclasses of Eulerian graphs: The first subclass consists of Eulerian *asymmetric* directed graphs¹ with maximum degree d and edge expansion at most α . Here we show a bound of $O(d^2\alpha^{-2})$ on the price of anarchy. The second subclass consists of unweighted d -regular Eulerian graphs, for which we obtain a bound of $d + 1$ on the price of anarchy.

Our Results: Modifying the Network: Finally, we consider an algorithmic problem within this framework of opinion formation. The question is the following: if we have the ability to modify the edges in the network (subject to certain constraints), how should we do this to reduce the social cost of the Nash equilibrium by as much as possible? This is a natural question both as a self-contained issue within the mathematical framework of opinion formation, and also in the context of applications: many social media sites overtly and algorithmically consider how to balance the mix of news content [1], [4], [16], [17] and also the mix of social content [3], [18] that they expose their users to, so as to optimize user engagement on the site.

We focus on three main variants on this question: when all edges must be added *to* a specific node (as in the case when a site can modify the amount of attention directed to a media source or celebrity); when all edges must be added *from* a specific node (as in the case when a particular media site tries to shift its location in the space of opinions by blending in content from others); and when edges can be added between any pair of nodes in the network (as in the case when a social networking site evaluates modifications to its feeds of content from one user to another [3], [18]).

Adding edges to reduce the social cost has an intuitive basis: it seems natural that exposing people to others with different opinions can reduce the extent of disagreement within the group. When one looks at the form of the social cost $c(y)$, however, there is something

¹An Eulerian *asymmetric* directed graph is an Eulerian graph that does not contain any pair of oppositely oriented edges (i, j) and (j, i) .

slightly counter-intuitive about the idea of adding edges to make things better: the social cost is a sum of quadratic terms, and by adding edges to G we are simply adding further quadratic terms to the cost. For this reason, in fact, adding edges to G can never improve the optimal social cost. But adding edges *can* improve the social cost of the Nash equilibrium, and sometimes by a significant amount — the point is that adding terms to the cost function shifts the equilibrium itself, which can sometimes more than offset the additional terms. For example, if we add a single edge from the center of the star in Figure 2 to one of the leaves, then the center will shift her opinion to $2/3$ in equilibrium, causing all the leaves to shift their opinions to $1/3$, and resulting in a $\Theta(n)$ improvement in the social cost. In this case, once the leader pays attention to even a single member of the group, the social cost improves dramatically.

In Section 4 we show that, in multiple variants, the problem of where to add edges to optimally reduce the social cost is NP-hard. But we obtain a set of positive results as well, including a $\frac{9}{4}$ -approximation algorithm when edges can be added between arbitrary pairs of nodes, and an algorithm to find the optimal amount of additional weight to add to a given edge.

2. UNDIRECTED GRAPHS

We first consider the case of undirected graphs and later handle the more general case of directed graphs. The main result in this section is a tight bound on the price of anarchy for the opinion-formation game in undirected graphs. After this, we discuss two slight extensions to the model: in the first, each player can put a different amount of weight on her internal opinion; and in the second, each player has multiple fixed opinions she listens to. We show that both models can be reduced to the basic form of the model that we study first.

For undirected graphs we can simplify the social cost to the following form:

$$c(z) = \sum_i (z_i - s_i)^2 + 2 \sum_{(i,j) \in E, i > j} w_{i,j} (z_i - z_j)^2.$$

We can write this concisely in matrix form, by using the *weighted Laplacian matrix* L of G . L is defined by setting $L_{i,i} = \sum_{j \in N(i)} w_{i,j}$ and $L_{i,j} = -w_{i,j}$. We can thus write the social cost as $c(z) = z^T A z + \|z - s\|^2$, where $A = 2L$. The optimal solution is the y minimizing $c(y)$. By taking derivatives, we see that the optimal solution satisfies $(A + I)y = s$. Since the Laplacian of a graph is a positive semidefinite matrix, it follows that $A + I$ is positive definite. Therefore, $(A + I)y = s$ has a unique solution: $y = (A + I)^{-1}s$.

In the Nash equilibrium each player chooses an opinion in order to minimize her cost; in terms of the derivatives of the cost functions, this implies that

$c'_i(x) = 0$ for all i . Thus, to find the opinions of the players in the Nash equilibrium we should solve the following system of equations: $\forall i (x_i - s_i) + \sum_{j \in N(i)} w_{i,j} (x_i - x_j) = 0$. Therefore in the Nash equilibrium each player holds an opinion which is a weighted average of her internal opinion and the Nash equilibrium opinions of all her neighbors. This can be succinctly written as $(L + I)x = (\frac{1}{2}A + I)x = s$. As before $\frac{1}{2}A + I$ is a positive definite matrix, and hence the unique Nash equilibrium is $x = (\frac{1}{2}A + I)^{-1}s$.

We now begin our discussion on the price of anarchy (PoA) of the opinion game — the ratio between the cost of the optimal solution and the cost of the Nash equilibrium.

Our main theorem is that the price of anarchy of the opinion game is at most $9/8$. Before proceeding to prove the theorem we present a simple upper bound of 2 on the PoA for undirected graphs. To see why this holds, note that the Nash equilibrium actually minimizes the function $z^T (\frac{1}{2}A)z + \|z - s\|^2$ (one can check that this function's partial derivatives are the system of equations defining the Nash equilibrium). This allows us to write the following bound on the PoA:

$$\begin{aligned} PoA = \frac{c(x)}{c(y)} &\leq \frac{2(x^T (\frac{1}{2}A)x + \|x - s\|^2)}{c(y)} \\ &\leq \frac{2(y^T (\frac{1}{2}A)y + \|y - s\|^2)}{c(y)} \\ &\leq \frac{2c(y)}{c(y)} = 2. \end{aligned}$$

We note that this bound holds only for the undirected case, as in the directed case the Nash equilibrium does not minimize $z^T (\frac{1}{2}A)z + \|z - s\|^2$.

We now state the main theorem of this section.

Theorem 2.1: For any graph G and any internal opinions vector s , the price of anarchy of the opinion game is at most $9/8$.

Proof: The crux of the proof is relating the price of anarchy of an instance to the eigenvalues of its Laplacian. Specifically, we characterize the graphs and internal opinion vectors with maximal PoA. In these worst-case instances at least one eigenvalue of the Laplacian is exactly 1, and the vector of internal opinions is a linear combination of the eigenvectors associated with the eigenvalues 1, plus a possible constant shift for each connected component. As a first step we consider two matrices B and C that arise by plugging the Nash equilibrium and optimal solution we previously computed into the cost function and applying simple

algebraic manipulations:

$$c(y) = s^T \underbrace{[(A+I)^{-1} - I]^2 + (A+I)^{-1}A(A+I)^{-1}}_B s$$

$$c(x) = s^T \underbrace{[(L+I)^{-1} - I]^2 + (L+I)^{-1}A(L+I)^{-1}}_C s.$$

The next step is to show that the matrices A, B, C are *simultaneously diagonalizable*: there exists an orthogonal matrix Q such that $A = Q\Lambda^A Q^T$, $B = Q\Lambda^B Q^T$ and $C = Q\Lambda^C Q^T$, where for a matrix M the notation Λ^M represents a diagonal matrix with the eigenvalues $\lambda_1^M, \dots, \lambda_n^M$ of M on the diagonal. We prove this in the full version, using basic facts about eigenvectors:

Lemma 2.2: A, B and C are simultaneously diagonalizable by a matrix Q whose columns are eigenvectors of A .

We can now express the PoA as a function of the eigenvalues of B and C . With $s' = Q^T s$ we have:

$$\begin{aligned} \text{PoA} &= \frac{c(x)}{c(y)} = \frac{s'^T C s'}{s'^T B s'} = \frac{s'^T Q \Lambda^C Q^T s'}{s'^T Q \Lambda^B Q^T s'} \\ &= \frac{s'^T \Lambda^C s'}{s'^T \Lambda^B s'} = \frac{\sum_{i=1}^n \lambda_i^C s_i'^2}{\sum_{i=1}^n \lambda_i^B s_i'^2} \leq \max_i \frac{\lambda_i^C}{\lambda_i^B} \end{aligned}$$

The final step of the proof consists of expressing λ_i^C and λ_i^B as functions of the eigenvalues of A (denoted by λ_i) and finding the value for λ_i maximizing the ratio between λ_i^C and λ_i^B .

Lemma 2.3: $\max_i \frac{\lambda_i^C}{\lambda_i^B} \leq 9/8$. The bound is tight if and only if there exists an i such that $\lambda_i = 2$.

Proof: Using basic facts about eigenvalues we get:

$$\begin{aligned} \lambda_i^B &= \left(1 - \frac{1}{\lambda_i + 1}\right)^2 + \frac{1}{\lambda_i + 1} \lambda_i \frac{1}{\lambda_i + 1} \\ &= \frac{\lambda_i^2}{(\lambda_i + 1)^2} + \frac{\lambda_i}{(\lambda_i + 1)^2} = \frac{\lambda_i^2 + \lambda_i}{(\lambda_i + 1)^2} = \frac{\lambda_i}{\lambda_i + 1} \\ \lambda_i^C &= \left(1 - \frac{1}{\frac{1}{2}\lambda_i + 1}\right)^2 + \frac{1}{\frac{1}{2}\lambda_i + 1} \lambda_i \frac{1}{\frac{1}{2}\lambda_i + 1} \\ &= \frac{\lambda_i^2}{(\lambda_i + 2)^2} + \frac{4\lambda_i}{(\lambda_i + 2)^2} = \frac{\lambda_i^2 + 4\lambda_i}{(\lambda_i + 2)^2}. \end{aligned}$$

We can now write $\lambda_i^C/\lambda_i^B = \phi(\lambda_i)$, where ϕ is a simple rational function:

$$\begin{aligned} \phi(\lambda) &= \frac{(\lambda^2 + 4\lambda)/(\lambda + 2)^2}{\lambda/(\lambda + 1)} = \frac{(\lambda^2 + 4\lambda)(\lambda + 1)}{(\lambda + 2)^2 \lambda} \\ &= \frac{(\lambda + 4)(\lambda + 1)}{(\lambda + 2)^2} = \frac{\lambda^2 + 5\lambda + 4}{\lambda^2 + 4\lambda + 4}. \end{aligned}$$

By taking the derivative of ϕ , we find that ϕ is maximized over all $\lambda \geq 0$ at $\lambda = 2$ and $\phi(2) = 9/8$.

The eigenvalues λ_i are all non-negative, so it is always true that $\max_i \phi(\lambda_i) \leq 9/8$. If 2 is an eigenvalue

of A (and hence 1 is an eigenvalue of the Laplacian) then there exists an internal opinions vector s for which the PoA is $9/8$. What is this opinion vector s ? To find it assume that the i th eigenvalue of the Laplacian equals 1. To get a PoA of $9/8$ we should choose $s'_i = 1$ and $\forall j \neq i$ $s'_j = 0$ in order to hit only λ_i . By definition $s' = Q^T s$, and hence $s = (Q^T)^{-1} s'$. Because Q is orthogonal, $Q^T = Q^{-1}$; thus, $s = Q s' = v_i$, where v_i is the eigenvector associated with λ_i . In fact, any linear combination of the eigenvectors associated with eigenvalues 0 and 1 where at least one of the eigenvectors of 1 has a coefficient different than 0 will obtain the maximal PoA. ■

With Lemma 2.3, we have completed the proof of Theorem 2.1. ■

Corollary 2.4: We can scale the weights of any graph to make its PoA be $9/8$. If α is the scaling factor for the weights, then the eigenvalues of the scaled A matrix are $\alpha \lambda_i$. Therefore by choosing $\alpha = \frac{2}{\lambda_i}$ for any eigenvalue other than 0 we get that there exists an internal opinions vector for which the PoA is $9/8$.

2.1. Arbitrary Node Weights and Players with Fixed Opinions

Our first extension is a model in which different people do not put the same weight on their internal opinion. In this extension, each node in the graph has a strictly positive weight w_i and the cost function is: $c(z) = \sum_i [w_i(z_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j}(z_i - z_j)^2]$. The

bound of $9/8$ on the PoA holds even in this model. To see this, let w be the vector of node weights and $d(w)$ be a diagonal matrix with the values of w on the diagonal. In terms of the scaled variables $\hat{z} = d(\sqrt{w})z$, $\hat{s} = d(\sqrt{w})s$ and the scaled matrix $\hat{A} = d(\sqrt{w})^{-1} A d(\sqrt{w})^{-1}$, the cost takes the same form as before: $c(\hat{z}) = \|\hat{z} - \hat{s}\|^2 + \hat{z}^T \hat{A} \hat{z}$. We have therefore proved:

Claim 2.5: The PoA of the game with arbitrary strictly positive node weights is bounded by $9/8$.

Next we show how to handle the case in which a subset of the players may have node weights of 0, which can equivalently be viewed as a set of players who have no internal opinion at all. We analyze this by first considering the case in which all non-zero node weights are the same; for this case we prove in the full version:

Claim 2.6: The PoA is bounded by $9/8$ if every player has either weight 1 or 0 on her internal opinion. By applying the change of variables from Claim 2.5 we can also handle non-zero arbitrary weights.

In the second model we present, some nodes have *fixed opinions*. In this model we partition the nodes into two sets A and B . Nodes in B are completely fixed in their opinion and are non-strategic, while nodes in A

have no internal opinion – they simply want to choose an opinion that minimizes the sum of their edges costs to their neighbors (which may include a mix of nodes in A and B). We can think of nodes in A as people forming their opinion and of nodes in B as news sources with a specific *fixed* orientation. We denote the fixed opinion of a node $j \in B$ by s_j . The social cost for this model is:

$$c(z) = \sum_{\substack{(i,j) \in E; \\ i \in A; j \in B}} (z_i - s_j)^2 + 2 \sum_{\substack{(i,j) \in E; \\ i,j \in A; i > j}} (z_i - z_j)^2.$$

Note that this clearly generalizes the original model, since we can construct a distinct node in B to represent each internal opinion. In the proof of Claim 2.7 in the full version we perform the reduction in the opposite direction, reducing this model to the basic model. To do this, we assign each node an internal opinion equal to the weighted average of the opinions of her fixed neighbors, and a weight equal to the sum of her fixed neighbors' weights. We then show that the PoA of the fixed opinion model is bounded by the PoA of the basic model and thus get:

Claim 2.7: The PoA of the fixed opinion model is at most $9/8$.

3. DIRECTED GRAPHS

We begin our discussion of directed graphs with an example showing that the price of anarchy can be unbounded even for graphs with bounded degrees. Our main result in this section is that we can nevertheless develop spectral methods extending those in Section 2 to find internal opinions that maximize the PoA for a given graph. Using this approach, we identify classes of directed graphs with good PoA bounds.

In the introduction we have seen that the PoA of an in-directed star can be unbounded. As a first question, we ask whether this is solely a consequence of the unbounded maximum in-degree of this graph, or whether it is possible to have an unbounded PoA for a graph with bounded degrees. Our next example shows that one can obtain a large PoA even when all degrees are bounded: we show that the PoA of a bounded degree tree can be $\Theta(n^c)$, where $c \leq 1$ is a constant depending on the in-degrees of the nodes in the tree.

Example 3.1: Let G be a 2^k -ary tree of depth $\log_{2^k} n$ in which the internal opinion of the root is 1 and the internal opinion of every other node is 0. All edges are directed toward the root. In the Nash equilibrium all nodes at layer i hold the same opinion, which is 2^{-i} . (The root is defined to be at layer 0.) The cost of a node at layer i is $2 \cdot 2^{-2i}$. Since there are 2^{ik} nodes at layer i , the total social cost at Nash equilibrium

is $\sum_{i=1}^{\log_{2^k} n} 2^{ik} 2^{1-2i} = 2 \sum_{i=1}^{\log_{2^k} n} 2^{(k-2)i}$. For $k > 2$ this

cost is $2^{k-1} \frac{(2^{k-2})^{\log_{2^k} n} - 1}{2^{k-2} - 1} = 2^{k-1} \frac{n^{\frac{k-2}{k}} - 1}{2^{k-2} - 1}$. The cost of the optimal solution is at most 1; in fact it is very close to 1, since in order to reduce the cost the root should hold an opinion of ϵ very close to 0, which makes the root's cost approximately 1. Therefore the PoA is $\Theta(n^{\frac{k-2}{k}})$. It is instructive to consider the PoA for extreme values of k . For $k = 2$, the PoA is $\Theta(\log n)$, while for $k = \log n$ we recover the in-directed star from the introduction where the PoA is $\Theta(n)$. For intermediate values of k , the PoA is $\Theta(n^c)$. For example, for $k = 3$ we get that the PoA is $\Theta(n^{\frac{1}{3}})$.

For directed graphs we do not consider the generalization to arbitrary node weights (along the lines of Section 2.1), noting instead that introducing node weights to directed graphs can have a severe effect on the PoA. That is, even in graphs containing only two nodes, introducing arbitrary node weights can make the PoA unbounded. For example, consider a graph with two nodes i and j . Node i has an internal opinion of 0 and a node weight of 1, while node j has an internal opinion of 1 and a node weight of ϵ . There is a directed edge (i, j) with weight 1. There exists a Nash equilibrium with a cost of $1/2$, but the social cost of the optimal solution is smaller than ϵ . To avoid this pathology, we restrict our attention to uniform node weights from now on.

3.1. The Price of Anarchy in a General Graph

For directed graphs the cost of the Nash equilibrium and the cost of the optimal solution are respectively $c(y) = s^T B s$ and $c(x) = s^T C s$, as before. But now, C has a slightly more complicated form since L is no longer a symmetric matrix. We first define the matrix A by setting $A_{i,j} = -w_{i,j} - w_{j,i}$ for $i \neq j$ and $A_{i,i} = \sum_{j \in N(i)} w_{i,j} + w_{j,i}$. The matrix A is the weighted Laplacian for an undirected graph where the weight on the undirected edge (i, j) is the sum of the weights in the directed graph for edges (i, j) and (j, i) . We then define $C = ((L + I)^{-1} - I)^T ((L + I)^{-1} - I) + (L + I)^{-T} A (L + I)^{-1}$. The price of anarchy, therefore, is $\frac{s^T C s}{s^T B s}$ as before. The primary distinction between the price of anarchy in the directed and undirected cases is that in the undirected case, B and C are both rational functions of A . In the directed case, no such simple relation exists between B and C , so that we cannot easily bound the generalized eigenvalues for the pair (and hence the price of anarchy) for arbitrary graphs. However, given a directed graph our main theorem shows that we can always find the vector of internal opinions s yielding the maximum PoA:

Theorem 3.2: Given a graph G it is possible to find the internal opinions vector s yielding the maximum PoA up to a precision of ϵ in polynomial time.

Proof: The total social cost is invariant under constant shifts in opinion. Therefore, without loss of generality, we restrict our attention to the space of opinion vectors with mean zero. Let us define a matrix $P \in \mathbb{R}^{n \times (n-1)}$ to have $P_{j,j} = 1$, $P_{j+1,j} = -1$, and $P_{i,j} = 0$ otherwise. The columns of P are a basis for the space of vectors with mean zero; that is, we can write any such vector as $s = P\hat{s}$ for some \hat{s} . We also define matrices $\bar{B} = P^T B P$ and $\bar{C} = P^T C P$, which are positive definite if the symmetrized graph is connected. The price of anarchy is then given by the generalized Rayleigh quotient $\rho_{\bar{C}, \bar{B}}(\hat{s}) = (\hat{s}^T \bar{C} \hat{s}) / (\hat{s}^T \bar{B} \hat{s})$. Stationary points of $\rho_{\bar{C}, \bar{B}}$ satisfy the generalized eigenvalue equation $(C - \rho_{\bar{C}, \bar{B}}(\hat{s})B)\hat{s} = 0$. In particular, the price of anarchy is the largest generalized eigenvalue, and the associated eigenvector \hat{s}_* corresponds to the maximizing choice of initial opinions.

The solution of generalized eigenvalue problems is a standard technique in numerical linear algebra, and there are good algorithms that run in polynomial time; see [12, §8.7]. In particular, because \bar{B} is symmetric and positive definite, we can use the Cholesky factorization $\bar{B} = R^T R$ to reduce the problem to the standard eigenvalue problem $(R^{-T} \bar{C} R^{-1} - \lambda I)(R\hat{s}) = 0$. ■

3.2. Upper Bounds for Classes of Graphs

In order to get good bounds on the PoA we restrict our attention to Eulerian graphs and pursue the following course of action: we begin by defining in Claim 3.3 a function $g(z)$ with the special property that its minimum value is the same as the cost of the Nash equilibrium. We next show in Claim 3.5 that by bounding $g(z)$ with a function of a specific structure we can get a bound on the PoA. Using this we present bounds for Eulerian bounded-degree asymmetric expanders, directed cycles, and the generalization of directed cycles to Eulerian d -regular graphs. As a first step, we use Schur complements to prove the following:

Claim 3.3: Let $g(z) = z^T M z + \|z - s\|^2$ with $M = (I - C)^{-1} - I$. If $(I - C)$ is nonsingular then for the Nash equilibrium x , we have $\min_z g(z) = c(x)$.

Proof: The social cost is a quadratic function of the expressed opinion vector and the internal opinion vector:

$$c(z) = z^T A z + \|z - s\|^2 = \begin{bmatrix} z \\ s \end{bmatrix}^T \begin{bmatrix} A + I & -I \\ -I & I \end{bmatrix} \begin{bmatrix} z \\ s \end{bmatrix}.$$

To compute the socially optimal vector, we minimize this quadratic form in z and s subject to constraints on s . This yields $c(y) = s^T B s$, where the matrix

$$B = ((A+I)^{-1} - I)^2 + (A+I)^{-1} A (A+I)^{-1} = I - (A+I)^{-1}$$

is a *Schur complement* in the larger system. Schur complements typically arise in partial elimination of variables from linear systems. In this case, we have eliminated the z variables in the stationary equations for a critical point in the extended quadratic form.

Now consider the Nash equilibrium. As we assume that 1 is not an eigenvalue of C , we can define

$$M = (I - C)^{-1} - I.$$

The matrix M is symmetric and positive semidefinite, with a null space consisting of the constant vectors. That is, we can see M as the Laplacian of a new graph. By design, $C = I - (M + I)^{-1}$, so we can mimic the construction above to express C as a Schur complement in a larger system. Thus, the social cost of the Nash equilibrium can be written

$$c(x) = \min_z \begin{bmatrix} z \\ s \end{bmatrix}^T \begin{bmatrix} M + I & -I \\ -I & I \end{bmatrix} \begin{bmatrix} z \\ s \end{bmatrix},$$

which is the optimal social cost in the new network. ■

We then complement the claim by showing that for Eulerian graphs $(I - C)$ is nonsingular and furthermore the matrix M has a nice structure:

Lemma 3.4: For Eulerian graphs $M = A + LL^T$.

Proof: We denote $\tilde{L} = L + I$ and $\tilde{A} = A + I$ then:

$$\begin{aligned} I - C &= I - (\tilde{L}^{-1} - I)^T (\tilde{L}^{-1} - I) - \tilde{L}^{-T} (\tilde{A} - I) \tilde{L}^{-1} \\ &= \tilde{L}^{-1} + \tilde{L}^{-T} - \tilde{L}^{-T} \tilde{A} \tilde{L}^{-1}. \end{aligned}$$

We use the fact that for Eulerian graphs $A = L + L^T$ which implies that $\tilde{A} = \tilde{L} + \tilde{L}^T - I$ to simplify $I - C$:

$$\begin{aligned} I - C &= \tilde{L}^{-1} + \tilde{L}^{-T} - \tilde{L}^{-T} (\tilde{L} + \tilde{L}^T - I) \tilde{L}^{-1} \\ &= \tilde{L}^{-T} \tilde{L}^{-1}. \end{aligned}$$

We have that $M = (L + I)(L + I)^T - I = A + LL^T$. Let us understand what the matrix LL^T looks like. On the diagonal we have $[LL^T]_{i,i} = d_i^2 + d_i$ where d_i is the degree of node i and off the diagonal $[LL^T]_{i,j} = d_i L_{j,i} + d_j L_{i,j} + \sum_{k \neq i,j} (L_{i,k} L_{j,k}) = d_i L_{j,i} + d_j L_{i,j} + |N(i) \cap N(j)|$. ■

Next, we show that by bounding $g(z)$ we can get bounds on the PoA:

Claim 3.5: Let G be connected, and let β be such that $g(z) \leq \beta(z^T A z) + \|z - s\|^2$. Then $\text{PoA} \leq \frac{\beta + \beta \lambda_2}{1 + \beta \lambda_2}$, where λ_2 is the second smallest eigenvalue of A .

Proof: Let \tilde{z} be the vector minimizing $g(z)$ and \tilde{y} the vector minimizing $\beta(\tilde{y}^T A \tilde{y}) + \|\tilde{y} - s\|^2$. Then we can derive the following bound on the price of anarchy:

$$\begin{aligned} \text{PoA}(G) &= \frac{c(x)}{c(y)} = \frac{g(\tilde{z})}{c(y)} \leq \frac{g(\tilde{y})}{c(y)} \\ &\leq \frac{\beta(\tilde{y}^T A \tilde{y}) + \|\tilde{y} - s\|^2}{(y^T A y) + \|y - s\|^2} = \frac{s^T C s}{s^T B s}, \end{aligned}$$

where C and B are defined similarly to the matrices in Theorem 2.1 and are simultaneously diagonalizable. If λ_i is an eigenvalue of A then $\lambda_i^B = \frac{\lambda_i}{1+\lambda_i}$ and $\lambda_i^C = \frac{\beta\lambda_i}{1+\beta\lambda_i}$. As before the maximum PoA is achieved when $\frac{\beta\lambda_i}{1+\beta\lambda_i} / \frac{\lambda_i}{1+\lambda_i} = \frac{\beta\lambda_i+\beta}{\beta\lambda_i+1}$ is maximized. Since all eigenvalues of A are positive we have that this is maximized for λ_2 . ■

Corollary 3.6: As an immediate corollary we have that for general Eulerian graphs the PoA is bounded by β . This holds even for weighted Eulerian graphs (each node's incoming weight is equal to its outgoing weight).

Furthermore, say that an Eulerian bounded degree *asymmetric* expander is an Eulerian graph that does not contain any pair of oppositely oriented edges (i, j) and (j, i) , and whose underlying undirected graph has maximum degree d and edge expansion α . We show:

Lemma 3.7: For an Eulerian bounded degree asymmetric expander with expansion α , the PoA is bounded by $O(d^2/\alpha^2)$.

Proof: For an asymmetric graph, the matrix A is the Laplacian of the underlying graph; this is why we require in the lemma that the graph is asymmetric.

If d is the maximum degree, then we have $\lambda_2 \leq \lambda_n \leq 2d$. We also have that $\lambda_2 \geq \alpha^2/2d$ [6]. We can now use this to bound the PoA in terms of the graph's expansion as follows:

$$\begin{aligned} \frac{\beta + \beta\lambda_2}{1 + \beta\lambda_2} &\leq \frac{\beta + \beta\lambda_2}{\beta\lambda_2} \leq \frac{1 + \lambda_2}{\lambda_2} \\ &\leq \frac{2d(1 + 2d)}{\alpha^2} = O(d^2/\alpha^2). \end{aligned}$$

■

This brings us to the following bound on the PoA, which is also a tight bound:

Claim 3.8: The PoA of a directed cycle is bounded by 2 and approaches 2 as the size of the cycle grows.

Proof: For a cycle we have $A = LL^T$; therefore $g(z) = 2(z^T Az) + \|z - s\|^2$, and hence the bound assumed in Claim 3.5 is actually a tight bound. In order to show that the PoA indeed approaches 2 we need to show that λ_2 approaches 0 as the size of the cycle grows. The fact that A is the Laplacian of an undirected cycle comes to our aid and provide us an exact formula for λ_2 : $\lambda_2 = 2(1 - \cos(\frac{2\pi}{n}))$ (where n is the size of the cycle), and this concludes the proof. ■

The bound can be generalized for Eulerian graphs:

Claim 3.9: The PoA of a d -regular Eulerian graph is bounded above by $d + 1$.

Proof: Similarly to directed cycles all we need to do is to show that $g(z) \leq (d + 1)(z^T Az) + \|z - s\|^2$ and the claim follows directly from Claim 3.5. We show

that $z^T(A + LL^T)z \leq (d + 1)z^T Az$:

$$\begin{aligned} z^T(LL^T)z &= (d^2 + d) \sum_i z_i^2 - 2d \sum_{(i,j) \in E} z_i z_j \\ &+ \sum_i \sum_{j>i} 2|N(i) \cap N(j)| z_i z_j. \end{aligned}$$

Notice that the last term can be bounded by $\sum_i d(d - 1)z_i^2$ using the fact that $x^2 + y^2 \geq 2xy$ and that in a d -regular graph two nodes can share at most $(d - 1)$ neighbors. Thus we get:

$$z^T(LL^T)z \leq d \left(2d \sum_i z_i^2 - 2 \sum_{(i,j) \in E} z_i z_j \right) = d(z^T Az).$$

■

For d -regular graphs we leave open the question of whether this is a tight bound or not. A more general open question which we leave open is whether there exists a Eulerian graph with PoA greater than 2.

4. ADDING EDGES TO THE GRAPH

In this section we consider the following class of problems: Given an unweighted graph G and a vector of internal opinions s , find edges E' to add to G so as to minimize the social cost of the Nash equilibrium. We begin with a general bound linking the possible improvement from adding edges to the price of anarchy. Let G be a graph (either undirected or directed). Denote by $c_G(z)$ the cost function and by x and y the Nash equilibrium and optimal solution respectively. Let G' be the graph constructed by adding edges to G . Then: $\frac{c_G(x)}{c_{G'}(x')} \leq \frac{c_G(x)}{c_{G'}(y')} \leq \frac{c_G(x)}{c_G(y)} = \text{PoA}(G)$. To see why this is the case, we first note that $c_{G'}(y') \leq c_{G'}(x')$ since the cost of the Nash equilibrium cannot be smaller than the optimal solution. Second, $c_G(y) \leq c_{G'}(y')$ simply because $c_{G'}(y)$ contains more terms than $c_G(y)$. Therefore we have proved the following claim:

Claim 4.1: Adding edges to a graph G can improve the cost of the Nash equilibrium by a multiplicative factor of at most the PoA of G .

We study three variants on the problem, discussed in the introduction. In all variants, we seek the “best” edges to add in order to minimize the social cost of the Nash equilibrium. The variants differ mainly in the types of edges we may add. First, we consider the case in which we can only add edges from a specific node w . Here we imagine that node w is a media source that therefore does not have any cost for holding an opinion, and so we will use a cost function that ignores the cost associated with it when computing the social cost. Hence, our goal is to find a set of nodes F such that adding edges from node w to all the nodes in F minimizes the cost of the Nash equilibrium while

ignoring the cost associated with w . By reducing the subset sum problem to this problem we show that:

Claim 4.2: Finding the best set of edges to add from a specific node w is NP-hard.

Next, we consider the case in which we can only add edges to a specific node; by reducing the minimum vertex cover problem to this problem we get that:

Claim 4.3: Finding the best set of edges to add to a specific node w is NP-hard.

The last case we consider is the most general one in which we can add any set of edges. For this case we leave open the question of the hardness of adding an unbounded set of edges. We do show that finding the best set of k arbitrary edges is NP-hard. This is done by a reduction from k -dense subgraph [9] :

Claim 4.4: Finding a best set of arbitrary k edges is NP-hard.

Finding approximation algorithms for all of the problems discussed in Claims 4.2, 4.3, and 4.4 is an interesting question. As a first step we offer a $\frac{9}{4}$ -approximation for the problem of optimally adding edges to a directed graph G — a problem whose hardness for exact optimization we do not know. The approximation algorithm works simply by including the reverse copy of every edge in G that is not already in G ; this produces a bi-directed graph G' .

Claim 4.5: $c_{G'}(x') \leq \frac{9}{4}c_G(y)$.

Proof: By Theorem 2.1 we have that $c_{G'}(x') \leq \frac{9}{8}c_{G'}(y')$. Also notice that in the worst case, in order to get from G to G' , we must double all the edges in G . Therefore $c_{G'}(y') \leq 2c_G(y)$. By combining the two we have that $c_{G'}(x') \leq \frac{9}{4}c_G(y)$. ■

For weighted graphs we can also include reverse copies of edges that do appear in G and therefore achieve an approximation ratio of 2 for analogous reasons.

4.1. Adding a Single Weighted Edge

We now consider how to optimally choose the weight to put on a single edge (i, j) , to minimize the cost of the Nash equilibrium. Suppose we add weight ρ to the edge (i, j) . The modified Laplacian is $L' = L + \rho e_i(e_i - e_j)^T$, where e_i is the i th vector in the standard basis. The modified Nash equilibrium is $x' = (L' + I)^{-1}s = ((L + I) + \rho e_i(e_i - e_j)^T)^{-1}s$.

Using the Sherman-Morrison formula for the rank-one update to an inverse [12, §2.1.3], we have

$$\begin{aligned} x' &= \left[(L + I)^{-1} - \frac{(L + I)^{-1} \rho e_i(e_i - e_j)^T (L + I)^{-1}}{1 + \rho(e_i - e_j)^T (L + I)^{-1} e_i} \right] s \\ &= x - v_i \left(\frac{\rho(x_i - x_j)}{1 + \rho(v_{i,i} - v_{i,j})} \right), \end{aligned}$$

where $v_i = (L + I)^{-1}e_i$ is the influence of s_i on the Nash opinions in the original graph. Therefore, v_i

gives the direction of change of the Nash equilibrium when the weight on (i, j) is increased: the equilibrium opinions all shift in the direction of v_i . We prove the following key properties of this influence vector v_i :

Lemma 4.6: The entries of $v_i = (L + I)^{-1}e_i$ lie in $[0, 1]$, and $v_{i,i}$ is the unique maximum entry.

Proof: The influence vector v_i is simply the Nash equilibrium for the internal opinion vector e_i . The Nash equilibrium is the limit of repeated averaging starting from the internal opinions, and the average of numbers in $[0, 1]$ is in $[0, 1]$. Thus the entries of v_i are in $[0, 1]$.

We show that $v_{i,i}$ is the maximal entry by contradiction. Suppose $v_{i,j}$ is maximal for some $j \neq i$. Because $L + I$ is nonsingular, v_i cannot be the zero vector, so $v_{i,j} > 0$. The equilibrium equations for j can be written

$$\begin{aligned} v_{i,j} &= \frac{\sum_{k \in N(j)} w_{j,k} v_{i,k}}{1 + \sum_{k \in N(j)} w_{j,k}} \\ &\leq \left(\frac{\sum_{k \in N(j)} w_{j,k}}{1 + \sum_{k \in N(j)} w_{j,k}} \right) \max_{k \in N(j)} v_{i,k} \leq \max_{k \in N(j)} v_{i,k} \end{aligned}$$

where the final inequality is strict if $v_{i,k} \neq 0$ for any $k \in N(j)$. But $v_{i,k} \neq 0$ for some $k \in N(j)$, since otherwise $v_{i,j}$ would be zero. Therefore, there must be some $k \in N(j)$ such that $v_{i,k} > v_{i,j}$, which contradicts the hypothesis that $v_{i,j}$ is maximal. ■

We now show how to choose the *optimal* weight ρ to add to edge (i, j) to best reduce the social cost of the Nash equilibrium.

Theorem 4.7: The optimal weight ρ to add to the edge (i, j) can be computed in polynomial time.

Proof: Note that

$$\begin{aligned} x'_i - x'_j &= (x_i - x_j) \left(1 - \frac{\rho(v_{i,i} - v_{i,j})}{1 + \rho(v_{i,i} - v_{i,j})} \right) \\ &= \frac{x_i - x_j}{1 + \rho(v_{i,i} - v_{i,j})}, \end{aligned}$$

and we can write the new Nash equilibrium as $x' = x - \phi v_i$, where

$$\phi = \frac{\rho(x_i - x_j)}{1 + \rho(v_{i,i} - v_{i,j})} = \rho(x_i - x_j) \frac{x'_i - x'_j}{x_i - x_j} = \rho(x'_i - x'_j).$$

For small values of ρ , we have that $\phi = \rho(x_i - x_j) + O(\rho^2)$; and as $\rho \rightarrow \infty$, we have that $\phi \rightarrow \phi_{\max} = (x_i - x_j)/(v_{i,i} - v_{i,j})$ and $x'_i - x'_j \rightarrow 0$. Thus, adding a small amount of weight to edge (i, j) moves the Nash equilibrium in the direction of the influence vector v_i proportional to the weight ρ and the discrepancy $x_i - x_j$; while adding larger amounts of weight moves the Nash equilibrium by a bounded amount in the direction of the influence vector v_i , with the asymptotic limit of large edge weight corresponding to the case when i and j have the same opinion.

What does adding a weighted edge between i and j do to the social cost at Nash equilibrium? In the modified graph, the social cost is

$$c'(z) = z^T A z + \rho(z_i - z_j)^2 + \|z - s\|^2.$$

At the new Nash equilibrium, we have

$$\begin{aligned} c'(x') &= x'^T A x' + \rho(x'_i - x'_j)^2 + \|x' - s\|^2 \\ &= x'^T A x' + \phi(x'_i - x'_j) + \|x' - s\|^2. \end{aligned}$$

Because x' is a linear function of ϕ , the above shows that $c'(x')$ is a quadratic function of ϕ , which we can simplify to $c'(x') = \alpha_{ij}\phi^2 - 2\beta_{ij}\phi + c(x)$, where

$$\begin{aligned} \alpha_{ij} &= v_i^T (A + I) v_i - (v_{i,i} - v_{i,j}) \\ \beta_{ij} &= v_i^T ((A + I)x - s) - \frac{1}{2}(x_i - x_j). \end{aligned}$$

The range of possible values for ϕ is between 0 (corresponding to $\rho = 0$) and ϕ_{\max} (corresponding to the limit as ρ goes to infinity). Subject to the constraints on the range of ϕ , the quadratic in ϕ is minimal either at 0, at ϕ_{\max} , or at β_{ij}/α_{ij} (assuming this point is between 0 and ϕ_{\max}). We can therefore determine the optimal weight for a single edge in polynomial time. ■

Note that the above computations also give us a simple formula for the gradient components γ_{ij} corresponding to differentiation with respect to w_{ij} :

$$\begin{aligned} \gamma_{ij} &\equiv \frac{d[c'(x')]}{d\rho} = \frac{d[c'(x')]}{d\phi} \frac{d\phi}{d\rho} = -2\beta_{ij}(x_i - x_j) \\ &= (x_i - x_j)^2 - 2(x_i - x_j)v_i^T((A + I)x - s). \end{aligned}$$

The residual vector $(A + I)x - s$ measures the extent to which x fails to satisfy the equation for the socially optimal opinion y . If this vector is large enough, and if the influence vector v_i is sufficiently well aligned with the residual, then adding weight to the (i, j) edge can decrease the social cost at Nash equilibrium. Thus, though computing a globally optimal choice of additional edge weights may be NP-hard, we can generally compute locally optimal edge additions via the method of steepest descent.

Acknowledgments: We thank Michael Macy for valuable discussion.

REFERENCES

- [1] D. Agarwal, B.-C. Chen, P. Elango, N. Motgi, S.-T. Park, R. Ramakrishnan, S. Roy, and J. Zachariah, "Online models for content optimization," in *Proc. NIPS*, 2008.
- [2] J. R. Alford, C. L. Funk, and J. R. Hibbing, "Are political orientations genetically transmitted?" *American Political Science Review*, vol. 99, no. 2, pp. 153–167, 2005.
- [3] L. Backstrom, E. Bakshy, J. Kleinberg, T. Lento, and I. Rosenn, "Center of attention: How Facebook users allocate attention across friends," in *Proc. 5th International Conference on Weblogs and Social Media*, 2011.
- [4] L. Backstrom, J. Kleinberg, R. Kumar, "Optimizing Web traffic via the media scheduling problem," *Proc. ACM SIGKDD Conf. Knowledge Disc. Data Mining*, 2009.
- [5] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," in *Proc. 7th Intl. Conf. on Computer Vision*, 1999.
- [6] F. R. K. Chung, *Spectral Graph Theory*. AMS 1997.
- [7] M. H. DeGroot, "Reaching a consensus," *J. American Statistical Association*, vol. 69, pp. 118–121, 1974.
- [8] P. M. DeMarzo, D. Vayanos, and J. Zweibel, "Persuasion bias, social influence, and unidimensional opinions," *Quarterly Journal of Economics*, vol. 118, no. 3, 2003.
- [9] U. Feige, D. Peleg, and G. Kortsarz, "The dense k -subgraph problem," *Algorithmica*, vol. 29, no. 3, 2001.
- [10] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *J. Math. Sociology*, vol. 15, no. 3-4, 1990.
- [11] B. Golub and M. O. Jackson, "Naive learning in social networks: Convergence, influence and the wisdom of crowds," *American Econ. J.: Microeconomics*, 2(2010).
- [12] G. H. Golub and C. F. V. Loan, *Matrix Computations*, 3rd ed. Johns Hopkins University Press, 1996.
- [13] M. O. Jackson, *Social and Economic Networks*. Princeton University Press, 2008.
- [14] J. M. Kleinberg and É. Tardos, "Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields," *Journal of the ACM*, vol. 49, no. 5, pp. 616–639, 2002.
- [15] D. Krackhardt, "A plunge into networks," *Science*, vol. 326, pp. 47–48, 2 October 2009.
- [16] S. A. Munson and P. Resnick, "Presenting diverse political opinions: how and how much," *Proc. 28th ACM Conf. Human Factors in Computing Sys. (CHI)*, 2010.
- [17] S. A. Munson, D. X. Zhou, and P. Resnick, "Sidelines: An algorithm for increasing diversity in news and opinion aggregators," in *Proc. 3rd Intl. Conf. Weblogs and Social Media*, 2009.
- [18] E. Sun, I. Rosenn, C. Marlow, and T. M. Lento, "Gesundheit! Modeling contagion through Facebook News Feed," *Proc. 3rd Intl. Conf. Weblogs and Social Media*, 2009.