

How Bad is Forming Your Own Opinion? *

David Bindel [†]

Jon Kleinberg [‡]

Sigal Oren [§]

June 10, 2014

Abstract

The question of how people form their opinion has fascinated economists and sociologists for long time. In many of the models, a group of people in a social network, each holding a numerical opinion, arrive at a shared opinion through repeated averaging with their neighbors in the network. Motivated by the observation that in reality consensus is rarely reached, we study a related sociological model in which individuals' intrinsic beliefs counterbalance the averaging process and yield a diversity of opinions.

We interpret the repeated averaging process as best-response dynamics in an underlying game with natural payoffs and its limit as an equilibrium. This allows us to study the cost of disagreement by comparing between the cost at equilibrium and the social optimum. We also consider a natural network design problem in this setting: which links can we add to the underlying network to reduce the cost of disagreement at equilibrium?

JEL Classification Numbers: C72, Z13

*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.

[†]Department of Computer Science, Cornell University, Ithaca NY 14853. Email: bindel@cs.cornell.edu.

[‡]Department of Computer Science, Cornell University, Ithaca NY 14853. Email: kleinber@cs.cornell.edu.

[§]Hebrew University and Microsoft Research, Israel. Email: sigalo@cs.huji.ac.il. The work was done while this author was at Cornell University.

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 in a social network can arrive at a shared opinion through a form of repeated averaging [1, 13, 16, 19]. This work builds on a basic model of DeGroot [12], 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. Now, in each time step node i updates her opinion to be a weighted average of her current opinion and 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, can only model 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 [21].

In this paper we study a model of opinion dynamics in which consensus is not reached in general, with the goal of quantifying 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 [15], which posits that each node i maintains a persistent *internal opinion* s_i . This internal opinion 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, in general, the presence of s_i as a constant in each iteration prevents repeated averaging from bringing 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 such as political orientations back to differences in education and background, and even to explore genetic bases for such patterns of variation [3].

Now, a line of work in sociology and cognitive psychology has argued that an individual's lack of agreement with his or her network neighbors creates *cognitive dissonance*, and individuals have an incentive to modify their behavior to reduce this dissonance (see [18] and the references therein).

Thus, when consensus is not reached, it becomes natural to ask about the cost of this lack of consensus, as a measure of the dissonance created.

What would be an appropriate way to measure this cost? As noted above, standard models of opinion evolution are based on repeated averaging, and so if we want to model a node’s opinion update as locally minimizing its own cost, then we should use a cost function that is optimized by averaging. This leads naturally to a quadratic cost function, since we can equivalently view a node’s averaging operation 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.

Given this view, we can think of repeated averaging as the trajectory of best-response dynamics in a one-shot, complete information 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).¹

Nash Equilibrium and Social Optimality in a Game of Opinion Formation

It was already observed in [15] that repeated averaging always converges. In Section 2 we repeat the convergence proof and show it actually converges 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 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. In fact we can view the problem of minimizing social cost for this game as a type of *metric labeling problem* [9, 20], albeit a polynomial-time solvable case of the problem 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 for this game is thus the *price of anarchy*, defined as the ratio between the cost of the Nash equilibrium and the cost of the optimal solution.

Our Results: Undirected Graphs

The model we have described can be used as stated in both undirected and directed graphs — the only difference is in whether i ’s neighbor set $N(i)$ represents the nodes with 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.

¹We note that subsequent to the appearance of the conference version of this work, Bhawalkar et al [8] extended the framework developed here to handle the class of strictly convex cost functions, rather than just quadratic functions. Chierichetti et al [10] extended the model in a different direction, to consider cost functions in which the quadratic terms in Equation (2) are replaced by distances in a discrete metric space.




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 $3/8$.

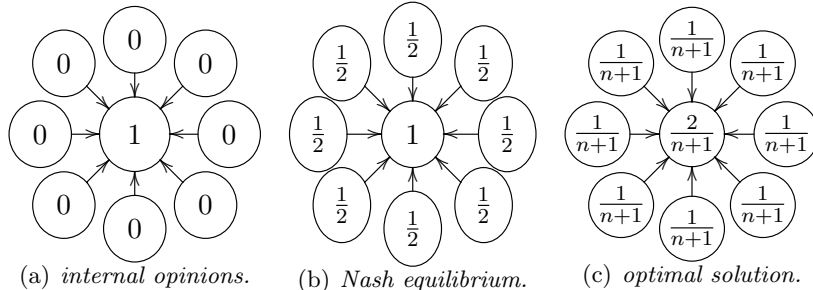


Figure 2: An example demonstrating that the price of anarchy of a directed graph can be unbounded.

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 with uniform edge weights 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 internal opinions vector 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 internal opinions vector 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 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 directed edges playing the role of undirected ones, but the range of possible behaviors in the model becomes very different. This is due 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 they maintain their internal opinion [19].

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, in contrast to the case of undirected graphs, the price of anarchy can be very large in some instances; 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 $\frac{1}{2}(n - 1)$, 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 4 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 framework for directed graphs to derive two bounds on the price of anarchy of Eulerian graphs: For Eulerian graphs with maximum degree Δ we obtain a bound of $\Delta + 1$ on the price of anarchy. For the subclass of Eulerian *antisymmetric* directed graphs² with maximum degree Δ and edge expansion α , we show a bound of $O(\Delta^2\alpha^{-2})$ 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 [2, 5, 22, 23] and also the mix of social content [4, 24] that they expose their users to, so as to optimize user engagement on the site.

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 slightly counter-intuitive about the idea of adding edges to improve the situation: 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

²An Eulerian *antisymmetric* directed graph is an Eulerian graph that does not contain any pair of oppositely oriented edges (i, j) and (j, i) .

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.

We focus on three main variants on this question. First, we consider a variant in which certain nodes correspond not to individuals but to media aggregators that simply blend content by linking to it, thus producing a hybrid opinion that can be viewed as the average of the opinions at the nodes that the aggregator links to. In this case, the question is how to add edges *from* the node corresponding to the aggregator so as to reduce the social cost by as much as possible. Second, we consider the case in which 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, which serves as the recipient of all the links. Finally, we consider the case in which edges can be added between any pair of nodes in the network; this corresponds to cases in which a social networking site evaluates modifications to its feeds of content from one user to another, since such modifications can alter which opinions an arbitrary user is exposed to. Indeed, modifications of this type have been the basis of recent experiments in opinion formation conducted by social networking sites [4, 6, 24].

In Section 5 we show that, in the previously discussed 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 Convergence and Nash Equilibrium Uniqueness

In this section, we show that the opinion game has a unique Nash equilibrium to which the repeated averaging process converges. Our proof makes use of some matrix notation which will also turn out to be useful for the rest of the paper. Given an edge-weighted graph, we write W for the weighted adjacency matrix, and let D be the diagonal matrix of node degrees $d_i = \sum_{j \neq i} w_{i,j}$. The *weighted graph Laplacian* is $L = D - W$. A useful fact about the Laplacian is that it is a *positive semidefinite* matrix. This implies that all of its eigenvalues are non-negative.

We are now ready to show that the Nash equilibrium is unique and is the limit of the repeated averaging process.

Claim 2.1 *The opinion game admits a unique Nash equilibrium.*

Proof: In the Nash equilibrium x each player chooses an opinion which minimizes her cost; in terms of the derivatives of the cost functions, this implies that $\frac{\partial}{\partial x_i} c_i(x) = 0$ for all i . Thus, to find the players' opinions 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. After some rearranging we get that $\forall i \sum_{j \in N(i)} w_{i,j}(x_i - x_j) + x_i = s_i$. This system of equations can be succinctly written as $(L + I)x = s$. Observe that, $L + I$ is a positive definite matrix as the Laplacian matrix is a positive semidefinite matrix and once we add the identity matrix to it all its eigenvalues are strictly greater than 0. Since a positive definite matrix is invertible, we have that the unique Nash equilibrium is $x = (L + I)^{-1}s$. \square

Claim 2.2 *The repeated averaging process defined by the Friedkin and Johnsen update rule (Equation 1) converges to the Nash equilibrium $x = (L + I)^{-1}s$.*

Proof: Let $z(t)$ be the opinions vector at time t . Let b be a “normalized” internal opinions vector defined as follows $b_i = \frac{s_i}{1 + \sum_{j \in N(i)} w_{i,j}}$. We also define a matrix R to be the “normalized” adjacency matrix of G : $R_{i,i} = 0$ and for $i \neq l$, $R_{i,l} = \frac{w_{i,l}}{1 + \sum_{j \in N(i)} w_{i,j}}$. With this notation in place the update rule in Equation 1 can now be written as $z(t) = F(z(t-1)) = Rz(t-1) + b$. If v^1 and v^2 are arbitrary vectors, then

$$\|F(v^1) - F(v^2)\|_\infty = \|R(v^1 - v^2)\|_\infty \leq \|R\|_\infty \|v^1 - v^2\|_\infty$$

where $\|v\|_\infty = \max_i |v_i|$ is the usual max norm and the associated matrix norm $\|R\|_\infty$ is given by

$$\|R\|_\infty = \max_i \sum_j |R_{i,j}| = \max_i \frac{\sum_{j \in N(i)} w_{i,j}}{1 + \sum_{j \in N(i)} w_{i,j}} < 1.$$

Thus, F is a contraction mapping in the max norm, and so the iteration converges to a unique fixed point. To compute this fixed point, we observe that $R = (D + I)^{-1}W$ and that $b = (D + I)^{-1}s$. Thus, we can alternatively write the repeated averaging process as $z(t+1) = (D + I)^{-1}(Wz(t) + s)$. After some rearranging we get that the fixed point of this process is indeed the Nash equilibrium: $x = (D + I - W)^{-1}s = (L + I)^{-1}s$. □

3 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 several fixed opinions she listens to instead of an internal opinion. We show that both models can be reduced to the basic form of the model which 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 write this concisely in matrix form as $c(z) = z^T A z + \|z - s\|^2$, where the matrix $A = 2L$ captures the tension on the edges. Recall that L is the weighted Laplacian of G and is defined by setting $L_{i,i} = \sum_{j \in N(i)} w_{i,j}$ and $L_{i,j} = -w_{i,j}$. The optimal solution is the y minimizing $c(\cdot)$. 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 comparison, as we showed in Claim 2.1, the Nash equilibrium is $x = (L + I)^{-1}s = (\frac{1}{2}A + I)^{-1}s$.

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 verify 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 $z^T(\frac{1}{2}A)z + \|z - s\|^2$ anymore.

We now state the main theorem of this section.

Theorem 3.1 *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:

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

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

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

Next, we 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.

Lemma 3.2 *A, B and C are simultaneously diagonalizable by a matrix Q whose columns are eigenvectors of A .*

Proof: It is a standard fact that any real symmetric matrix M can be diagonalized by an orthogonal matrix Q such that $M = Q\Lambda^M Q^T$. Q 's columns are eigenvectors of M which are orthogonal to each other and have a norm of one. Thus in order to show that A , B and C can be diagonalized with the same matrix Q it is enough to show that all three are symmetric and have the same eigenvectors. For this we use the following basic fact:

If λ^N is an eigenvalue of N , λ^M is an eigenvalue of M and w is an eigenvector of both then:

1. $\frac{1}{\lambda^M}$ is an eigenvalue of M^{-1} and w is an eigenvector of M^{-1} .
2. $\lambda^N + \lambda^M$ is an eigenvalue of $N + M$ and w is an eigenvector of $N + M$.
3. $\lambda^N \cdot \lambda^M$ is an eigenvalue of NM and w is an eigenvector of NM .

From this we can show that any eigenvector of A is also an eigenvector of B and C . Recall that A is a symmetric matrix, thus, it has n orthogonal eigenvectors which implies that A, B and C are all symmetric and share the same basis of eigenvectors. Therefore A, B and C are simultaneously diagonalizable. \square

We can now express the PoA as a function of the eigenvalues of B and C . By defining $\hat{s} = Q^T s$ we have:

$$\begin{aligned} PoA &= \frac{c(x)}{c(y)} = \frac{s^T C s}{s^T B s} = \frac{s^T Q \Lambda^C Q^T s}{s^T Q \Lambda^B Q^T s} \\ &= \frac{\hat{s}^T \Lambda^C \hat{s}}{\hat{s}^T \Lambda^B \hat{s}} = \frac{\sum_{i=1}^n \lambda_i^C \hat{s}_i^2}{\sum_{i=1}^n \lambda_i^B \hat{s}_i^2} \leq \max_i \frac{\lambda_i^C}{\lambda_i^B} \end{aligned}$$

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 3.3 $\max_i \frac{\lambda_i^C}{\lambda_i^B} \leq 9/8$. *The bound is tight if and only if there exists an i such that $\lambda_i = 2$.*

Proof: Using the basic facts about eigenvalues which were mentioned in the proof of Lemma 3.2, we get:

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

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

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

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 the internal opinions vector maximizing the PoA? Rewriting our expression from above, we have

$$PoA = \frac{\sum_{i=1}^n \hat{s}_i^2 \lambda_i^B \phi(\lambda_i)}{\sum_{i=1}^n \hat{s}_i^2 \lambda_i^B},$$

i.e. the price of anarchy is a weighted average of the values $\phi(\lambda_i)$, where the weights are given by $\hat{s}_i^2 \lambda_i^B$. The only way to achieve the maximum value is if the only nonzero weights are on eigenvalues maximizing $\phi(\lambda)$. Because λ_i^B is positive whenever λ_i is positive, this means that to achieve a PoA of $9/8$, \hat{s}_i^2 can only be nonzero if $\lambda_i = 2$ or $\lambda_i = 0$. Recall that $s = Q\hat{s}$, where the columns of Q are the eigenvectors of A . Thus, any internal opinion vector that is an eigenvector of A with eigenvalue 2 plus some null vector of A will achieve the maximal price of anarchy. In particular, since the all-ones vector is always an eigenvector of the weighted Laplacian matrix (and thus A) we have that there exists an internal opinions vector where all the opinions are positive for which the maximal PoA is achieved. \square

With Lemma 3.3, we have completed the proof of Theorem 3.1. \square

Corollary 3.4 *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 $9/8$.*

3.1 Arbitrary Node Weights and Players with Fixed Opinions

Our first extension is a model in which different people put different weights on their internal opinion. In this extension, each node in the graph has a strictly positive weight γ_i and the cost function is

$$c^1(z) = \sum_i \left(\gamma_i (z_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j} (z_i - z_j)^2 \right) = (z - s)^T \Gamma (z - s) + z^T A z,$$

where Γ is the diagonal matrix of node weights γ_i and $A = 2L$ is defined as in the previous section. In the next claim we show that the bound of $9/8$ on the PoA holds even in this model:

Claim 3.5 *The PoA of the game with arbitrary strictly positive node weights is bounded by $9/8$.*

Proof: We define the scaled variables $\hat{z} = \Gamma^{1/2} z$, $\hat{s} = \Gamma^{1/2} s$ and the scaled matrices $\hat{A} = \Gamma^{-1/2} A \Gamma^{-1/2}$ and $\hat{L} = \Gamma^{-1/2} L \Gamma^{-1/2} = \hat{A}/2$. Now, in terms of the scaled variables, we have

$$\begin{aligned} c^1(z) &= (z - s)^T \Gamma (z - s) + z^T A z \\ &= (\hat{z} - \hat{s})^T (\hat{z} - \hat{s}) + \hat{z}^T \hat{A} \hat{z} \\ &= \|\hat{z} - \hat{s}\|^2 + \hat{z}^T \hat{A} \hat{z}, \end{aligned}$$

and the Nash equation $(L + \Gamma)x = \Gamma s$ can similarly be multiplied by $\Gamma^{-1/2}$ and rewritten in terms of the scaled variables as $(\hat{L} + I)\hat{x} = \hat{s}$. Thus, in terms of the scaled variables the problem takes exactly the same form as in the previous section, and the argument of Theorem 3.1 applies. \square

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 all³.

³Note that the repeated averaging process defined by the Friedkin and Johnsen update rule converges as long as at least one player in each connected component has nonzero weight. This is due to the fact that the Friedkin and Johnsen update rule is equivalent to the update rule for Jacobi iteration for the linear system $(L + I)x = s$, and Claim 2.2 is a restatement of the theorem that Jacobi iteration converges for strictly diagonally dominant matrices. Jacobi iteration also converges for irreducibly diagonally dominant matrices [25, Theorem 3.4, page 73].

Lemma 3.6 *The PoA of the game with non-negative node weights is bounded by 9/8.*

Proof: We assume without loss of generality that every connected component of G includes at least one player i with weight $\gamma_i > 0$. Observe that otherwise, the cost associated with this connected component both in the Nash equilibrium and in the optimal solution is 0 and hence we can ignore this component. We begin by showing the matrix $L + \Gamma$ is positive definite, and thus nonsingular. Both L and Γ are positive semidefinite, so we only need to show that their sum is positive definite. If $z^T L z = 0$, then z is constant over each connected component of G ; and if $z^T \Gamma z = 0$, then z is zero for at least one node in each component. Therefore, $z^T L z = z^T \Gamma z = 0$ if and only if $z = 0$. Because $A=2L$, the same argument shows that $A + \Gamma$ is nonsingular. Thus, the Nash equilibrium equations $(L + \Gamma)x = \Gamma s$ and the social optimality equation $(A + \Gamma)y = \Gamma s$ each have a unique solution.

To show the PoA is bounded by 9/8, consider a modified problem with weights $\gamma_i + \epsilon$ for $\epsilon \geq 0$. Because $L + \Gamma + \epsilon I$ and $A + \Gamma + \epsilon I$ are nonsingular for all $\epsilon \geq 0$, the Nash equilibrium $x^\epsilon = (L + \Gamma + \epsilon I)^{-1}(\Gamma + \epsilon I)s$ and the social optimum $y^\epsilon = (A + \Gamma + \epsilon I)^{-1}(\Gamma + \epsilon I)s$ are both continuous functions of ϵ . The modified social cost $c^{1,\epsilon}(z)$ is a continuous function of both ϵ and z . Therefore,

$$PoA^\epsilon = \frac{c^{1,\epsilon}(x^\epsilon)}{c^{1,\epsilon}(y^\epsilon)}$$

is a continuous function in ϵ , provided that $c^{1,\epsilon}(y^\epsilon) \neq 0$. Thus, there are two cases:

1. At least one connected component of G includes nodes i and j with γ_i and γ_j both positive and $s_i \neq s_j$. In this case $c^{1,\epsilon}(y^\epsilon) \neq 0$ for any $\epsilon > 0$, and PoA^ϵ is a continuous function of ϵ for any $\epsilon \geq 0$. Because $PoA^\epsilon \leq 9/8$ for $\epsilon > 0$, continuity implies $PoA^0 \leq 9/8$.
2. No connected component of G includes nodes i and j with γ_i and γ_j both positive and $s_i \neq s_j$. In this case, $x^\epsilon = y^\epsilon$ is constant on each connected component, and $c^{1,\epsilon}(x^\epsilon) = c^{1,\epsilon}(y^\epsilon) = 0$ for any $\epsilon \geq 0$. In this case, we define the price of anarchy to be 1, which is bounded by 9/8.

□

In the second model we present, some nodes have *fixed opinions* and others do not have an internal opinion at all. 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 their disagreement with 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 cost for player $i \in A$ in this model is

$$c_i^2(z) = \sum_{j:(i,j) \in E_{AB}} w_{i,j}(z_i - s_j)^2 + \sum_{j:(i,j) \in E_{AA}} w_{i,j}(z_i - z_j)^2,$$

where E_{AB} and E_{AA} denote the edges between A and B and between A and A , respectively. Note that this clearly generalizes the original model, since we can construct a distinct node in B to represent each internal opinion. Next, 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:

Proposition 3.7 *The PoA of the fixed opinion model is at most 9/8.*

Proof: We assume without loss of generality that any component of the subgraph G' on nodes in A has an edge to some node in B . This is a valid assumption as the cost associated with any component that do not have such an edge is 0 both in the Nash equilibrium and in the optimal solution. Hence, we can ignore such components. For each player $i \in A$ with edges into B , define the total edge weight γ_i and the weighted mean s_i and variance ν_i of the fixed opinions of neighbors as

$$\begin{aligned}\gamma_i &= \sum_{j:(i,j) \in E_{AB}} w_{i,j} \\ s_i &= \frac{1}{\gamma_i} \sum_{j:(i,j) \in E_{AB}} w_{i,j} s_j \\ \nu_i &= \frac{1}{\gamma_i} \sum_{j:(i,j) \in E_{AB}} w_{i,j} (s_j^2 - s_i^2).\end{aligned}$$

For convenience, let γ_i , s_i , and ν_i be zero for nodes with no edges into B . Note that

$$\begin{aligned}\sum_{j:(i,j) \in E_{AB}} w_{i,j} (z_i - s_j)^2 &= \left(\sum_{j:(i,j) \in E_{AB}} w_{i,j} \right) z_i^2 - 2 \left(\sum_{j:(i,j) \in E_{AB}} w_{i,j} s_j \right) z_i + \left(\sum_{j:(i,j) \in E_{AB}} w_{i,j} s_j^2 \right) \\ &= \gamma_i z_i^2 - 2\gamma_i s_i z_i + \gamma_i (\nu_i + s_i^2) \\ &= \gamma_i (z_i - s_i)^2 + \gamma_i \nu_i.\end{aligned}$$

Thus, the cost for a player $i \in A$ in the fixed opinion game is

$$c_i^2(z) = \gamma_i \nu_i + \gamma_i (z_i - s_i)^2 + \sum_{j:(i,j) \in E_{AA}} w_{i,j} (z_i - z_j)^2 = c_i^1(z) + \gamma_i \nu_i,$$

where $c^1(z)$ is the cost for the game played on the subgraph G' over the nodes in A with node weights γ_i and intrinsic opinions s_i . Similarly, the social cost in the fixed opinion game is

$$c^2(z) = c^1(z) + \nu,$$

where $\nu = \sum_i \gamma_i \nu_i$. Because the variance costs are independent of z , we have that $\partial c_i^2 / \partial z_j = \partial c_i^1 / \partial z_j$; thus, the equations for the Nash opinion vector and the optimal opinion vector are the same as in the previous model with arbitrary weights. From Lemma 3.6, we know $c^1(x) \leq \frac{9}{8} c^1(y)$; therefore,

$$c^2(x) = c^1(x) + \nu \leq \frac{9}{8} c^1(y) + \nu \leq \frac{9}{8} c^1(y) + \frac{9}{8} \nu = \frac{9}{8} c^2(y).$$

Hence, the price of anarchy is again bounded by 9/8. □

4 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 3 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 4.1 *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 3.1), 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. The cost of the Nash equilibrium is $1/2$, but the social cost of the optimal solution is smaller than ϵ . Thus, from now on we restrict our attention to uniform node weights.

4.1 The Price of Anarchy in a General Graph

For directed graphs we can define matrices B and C , similarly to their definition for undirected graphs, such that the cost of the optimal solution and the cost of the Nash equilibrium are respectively $c(y) = s^T B s$ and $c(x) = s^T C s$. Recall that the matrix A is used in the social cost function to capture the cost associated with the edges of the graph (disagreement between neighbors). We define it for directed graphs by setting $A_{i,j} = -w_{i,j} - w_{j,i}$ for $i \neq j$ and $A_{i,i} = \sum_{j \in N(i)} w_{i,j} + \sum_{\{j | i \in N(j)\}} w_{j,i}$. The matrix A can be also interpreted as 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) . Note that A is no longer a linear function of L , which is what makes analyzing the PoA of directed graphs more challenging. Recall that the matrices B and C are defined as follows:

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

$$C = (L + I)^{-T} (A + L^T L) (L + I)^{-1}.$$

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 4.2 *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 for internal opinion vector s is $\frac{\hat{s}^T \bar{C} \hat{s}}{\hat{s}^T \bar{B} \hat{s}}$ which is also known as the generalized Rayleigh quotient $\rho_{\bar{C}, \bar{B}}(\hat{s})$. To compute the maximum value that the PoA can take we observe that the stationary points of $\rho_{\bar{C}, \bar{B}}(\cdot)$ satisfy the generalized eigenvalue equation $(\bar{C} - \rho_{\bar{C}, \bar{B}}(\hat{s})\bar{B})\hat{s} = 0$. In particular, the maximal price of anarchy is the largest generalized eigenvalue, and the associated eigenvector \hat{s}_* corresponds to the maximizing choice of internal 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 [17, §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 $(R^{-T} \bar{C} R^{-1} - \lambda I)(R\hat{s}) = 0$. \square

4.2 Upper Bounds for Classes of Graphs

Our goal in this section is rather simple: we would like to find families of graphs for which we can bound the price of anarchy. The main tool we use is bounding the cost of the Nash equilibrium by a function of a simple structure. By using a function that has a similar structure to the social cost function we are able to frame the bound as a generalized eigenvalue problem that can be solved using techniques similar to the ones that were used in proving Theorem 3.1.

Proposition 4.3 *Let \mathcal{G} be a graph family for which there exists a β such that for any $G \in \mathcal{G}$ and any internal opinion vector s , we have $c(x) \leq \min_z \tilde{c}(z)$, where $\tilde{c}(z) = \beta(z^T A z) + \|z - s\|^2$. Then, $\forall G \in \mathcal{G}$ and internal opinion vector s , $\text{PoA}(G) \leq \frac{\beta + \beta \lambda_2}{1 + \beta \lambda_2}$, where λ_2 is the second smallest eigenvalue of A .*

Proof: Let $\tilde{y} = (\beta A + I)^{-1} s$ be the vector minimizing $\tilde{c}(\cdot)$. We can derive the following bound on the price of anarchy:

$$\text{PoA}(G) = \frac{c(x)}{c(y)} \leq \frac{\tilde{c}(\tilde{y})}{c(y)} = \frac{s^T \tilde{C} s}{s^T B s},$$

where \tilde{C} and B are defined similarly to the matrices in Theorem 3.1:

$$\begin{aligned} B &= I - (A + I)^{-1} \\ \tilde{C} &= I - (\beta A + I)^{-1} \end{aligned}$$

and are simultaneously diagonalizable. If λ_i is an eigenvalue of A then $\lambda_i^B = \frac{\lambda_i}{1+\lambda_i}$ and $\lambda_i^{\tilde{C}} = \frac{\beta\lambda_i}{1+\beta\lambda_i}$. As before, the maximum PoA is achieved when $\lambda_i^{\tilde{C}}/\lambda_i^B = \frac{\beta\lambda_i}{1+\beta\lambda_i} / \frac{\lambda_i}{1+\lambda_i} = \frac{\beta\lambda_i + \beta}{\beta\lambda_i + 1}$ is maximized. The maximum here is taken over all eigenvalues different than 0 as we know that the PoA for the internal opinions vector associated with eigenvalue 0 (which is a constant vector) is 1. Therefore, for a connected graph the maximizing eigenvalue is λ_2 . \square

An immediate corollary is that if there exists a β as in Proposition 4.3 then the PoA is bounded by this β .

We say that an unweighted bounded degree *antisymmetric* expander is an unweighted directed graph that does not contain any pair of oppositely oriented edges (i, j) and (j, i) , and whose symmetrized graph has maximum degree Δ and edge expansion α . We show:

Claim 4.4 *Let $\mathcal{G}_e \subseteq \mathcal{G}$ be a graph family consisting of unweighted bounded degree antisymmetric expanders for which the β defined in Proposition 4.3 exists. The PoA of $G \in \mathcal{G}_e$ is bounded by $O(\Delta^2/\alpha^2)$.*

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

If Δ is the maximum degree, then we have $\lambda_2 \leq \lambda_n \leq \Delta$. We also have that $\lambda_2 \geq \alpha^2/2\Delta$ [11]. We can now use this to bound the expression we got in Proposition 4.3 for the PoA in terms of the graph's expansion as follows:

$$\frac{\beta + \beta\lambda_2}{1 + \beta\lambda_2} \leq \frac{\beta + \beta\lambda_2}{\beta\lambda_2} \leq \frac{1 + \lambda_2}{\lambda_2} \leq \frac{2\Delta(1 + \Delta)}{\alpha^2} = O(\Delta^2/\alpha^2).$$

\square

The next natural question is for which graph families such a β exists. Intuitively, such a β exists whenever the cost of the Nash equilibrium is smaller than the cost of the best consensus — that is, the optimal solution restricted to opinion vectors in which all players hold the same opinion (constant vectors). This is true since the function $\beta(z^T Az) + \|z - s\|^2$ is the social cost function of a network in which the weights of all edges have been multiplied by β . However using this intuition for finding graph families for which β exists is difficult and furthermore does not help in computing the value of β (or a bound on it). Hence, we take a different approach. In Lemma 4.5, we introduce an intermediate function $g(\cdot)$ with the special property that its minimum value is the same as the cost of the Nash equilibrium. By showing that there exists a β such that $g(z) \leq \beta z^T Az + \|z - s\|^2$ we are able to present bounds for Eulerian bounded-degree graphs and additional bounds for Eulerian bounded-degree antisymmetric expanders. As a first step, we prove the following:

Lemma 4.5 *For unweighted Eulerian graphs, the social cost at Nash equilibrium is $c(x) = \min_z g(z)$, where $g(z) = z^T M z + \|z - s\|^2$ and $M = A + LL^T$.*

Proof: In the Eulerian case we have $A = L + L^T$, and we use this to simplify the expression

$$C = (L + I)^{-T} (A + L^T L) (L + I)^{-1}.$$

We first note that

$$A + L^T L = L + L^T + L^T L = (L + I)^T (L + I) - I,$$

then substitute to find

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

Where the last transition was based in the fact that

$$(L + I)(L + I)^T - I = L + L^T + LL^T = A + LL^T = M$$

Because M is positive semidefinite, by taking derivatives we have that the unique minimizer of $g(z)$ is $\hat{y} = (M + I)^{-1}s$; proceeding as in the derivation for the cost at optimality, we find the minimum value achieved is

$$g(\hat{y}) = \hat{y}^T(M + I)\hat{y} - 2s^T\hat{y} + s^Ts = s^T [I - (M + I)^{-1}] s = s^T C s = c(x).$$

□

Recall that Δ is the maximum degree of an Eulerian graph. We are now ready to prove the following proposition:

Proposition 4.6 *For unweighted bounded degree Eulerian graphs $c(x) \leq \min_z(\Delta + 1)(z^T A z) + \|z - s\|^2$.*

Proof: By Lemma 4.5 we have that for Eulerian graphs $c(x) = \min_z g(z) = \min_z z^T(A + LL^T)z + \|z - s\|^2$. What remains to show is that for $\beta = \Delta + 1$ it holds that $g(z) \leq \beta z^T A z + \|z - s\|^2$. After some rearranging this boils down to showing that the following holds: $z^T LL^T z \leq (\beta - 1)z^T A z$.

Note that A is the Laplacian for a symmetrized version of the graph; assuming this graph is connected (since otherwise we can work separately in each component), this means A has one zero eigenvalue corresponding to the constant vectors, and is positive definite on the space orthogonal to the constant vector. Similarly, LL^T has a zero eigenvalue corresponding to the constant vectors, and is at least positive semi-definite on the space orthogonal to the constant vectors. Since A is positive definite on the space of non-constant vectors, the smallest possible β can be computed via the solution of a generalized eigenvalue problem

$$\beta = 1 + \max_{z \neq \alpha e} \frac{z^T LL^T z}{z^T A z},$$

where e denotes the all-ones vector. In the case of an unweighted graph, one get a bound via norm inequalities. Using the fact that the graph is Eulerian, L^T is also a graph Laplacian, and we can write

$$(L^T z)_i = \sum_{j=1}^n w_{j,i}(z_i - z_j),$$

so

$$z^T LL^T z = \sum_{i=1}^n \left(\sum_{j=1}^n w_{j,i}(z_i - z_j) \right)^2.$$

Similarly, we expand the quadratic form $z^T Az$ into

$$z^T Az = \sum_{i < j} (w_{i,j} + w_{j,i})(z_i - z_j)^2 = \sum_{i=1}^n \left(\sum_{j=1}^n w_{j,i}(z_i - z_j)^2 \right).$$

Now, recall that in general $\left(\sum_{j=1}^d x_j\right)^2 \leq d \sum_{j=1}^d x_j^2$, which means that in the unweighted case $\left(\sum_{j=1}^n w_{j,i}(z_i - z_j)\right)^2 \leq d_i \left(\sum_{j=1}^n w_{j,i}(z_i - z_j)^2\right)$, where $d_i = \sum_j w_{j,i}$ is the in-degree of node i (which is the same as the out-degree). Therefore,

$$\frac{z^T LL^T z}{z^T Az} \leq \frac{\sum_{i=1}^n d_i \sum_{j=1}^n w_{j,i}(z_i - z_j)^2}{\sum_{i=1}^n \sum_{j=1}^n w_{j,i}(z_i - z_j)^2} \leq \max_i d_i = \Delta.$$

So for a general Eulerian graph, $\beta \leq 1 + \Delta$. □

We observe that for a cycle the bound of 2 on the price of anarchy is actually tight:

Observation 4.7 *The PoA of a directed cycle is bounded by 2 and approaches 2 as the size of the cycle grows.*

Proof: For a cycle it is the case that $A = LL^T$; therefore $g(z) = 2(z^T Az) + \|z - s\|^2$, and hence the bound assumed in Proposition 4.3 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 ([11]), and this concludes the proof. □

For general Eulerian graphs we leave open the question of whether the bound of $\Delta + 1$ is a tight bound or not. Indeed, it is an intriguing open question whether there exists a Eulerian graph with PoA greater than 2.

5 Adding Edges to the Graph

The next thing we consider is 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 $c_G(\cdot)$ 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 $c_{G'}(\cdot)$ contains more terms than $c_G(\cdot)$. Therefore we have proved the following proposition:

Proposition 5.1 *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.

Adding edges from a specific node

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 exhibited by w . By reducing the subset sum problem to this problem we show that:

Proposition 5.2 *Finding the best set of edges to add from a specific node w is NP-hard.*

Proof: Denote by $G + F$ the graph constructed by adding to G edges from w to all nodes in F . Our goal is to find a set F minimizing $\tilde{c}_{G+F}(x)$, where x is a Nash equilibrium in the graph $G + F$ and \tilde{c} denotes the total cost of all nodes in x except for node w . We show that finding this set is NP-hard by reducing the subset sum problem to this problem. Recall that in the subset sum problem we are given a set of positive integers a_1, \dots, a_n and a number t . We would like to know if there exists any subset S such that $\sum_{j \in S} a_j = t$. Given an instance of the subset problem, we reduce it to the following instance of the opinion game. The instance contains an in-directed star with n peripheral nodes that have an internal opinion of 0 and a center node w which has an internal opinion of 1 and n isolated nodes that have internal opinions of $-\frac{a_i}{t}$. This construction is illustrated in Figure 3.




Figure 3: An illustration of the instance structured for Proposition 5.2.

Lemma 5.3 *For the graph G and the vector of internal opinions s defined above, there exists a set F such that $\tilde{c}_{G+F}(x) = 0$ if and only if the answer to the subset problem is yes.*

Proof: As seen in the introduction, in the Nash equilibrium each one of the peripheral nodes holds an opinion of $\frac{1}{2}x_w$. Node w hold an opinion of $x_w = \frac{1 + \sum_{j \in F} s_j}{1 + |F|}$. As we assume that w does not incur any cost, the cost of the Nash equilibrium in $G + F$ is just the cost of the n peripheral nodes:

$$\tilde{c}_{G+F}(x) = n \left(\left(\frac{1}{2}x_w - 0 \right)^2 + \left(x_w - \frac{1}{2}x_w \right)^2 \right) = 2n \left(\frac{1 + \sum_{j \in F} s_j}{2(1 + |F|)} \right)^2.$$

Clearly the cost is nonnegative as it is a sum of quadratic terms; moreover it equals 0 if and only if $\sum_{j \in F} s_j = -1$. Defining $F' = \{j \in F | s_j < 0\}$, we have $\sum_{j \in F'} s_j = -1$. By the reduction we have

that $\sum_{j \in F'} -\frac{a_j}{t} = -1$; if we multiply by $-t$ we get that $\sum_{j \in F'} a_j = t$ implying that there exists a solution to the subset sum problem. \square \square

Adding edges to a specific node

Next, we consider the case in which we can only add edges to a specific node. We can imagine again that node w is a media source; in this case, however, our goal is to find the best set of people to expose to this media source. By reducing the minimum vertex cover problem to this problem we show that:

Proposition 5.4 *Finding the best set of edges to add to a specific node w is NP-hard.*

Proof: Given an instance of the minimum vertex cover problem, consisting of an undirected graph $G' = (V', E')$, we construct an instance of the opinions game as follows:

- For each edge $(i, j) \in E'$ we create a vertex $v_{i,j}$ with internal opinion 1.
- For every $v_{i,j}$ we create an in-directed star with 24 peripheral nodes that have an internal opinion of 0. We later refer to node $v_{i,j}$ and all the nodes directed to it as $v_{i,j}$'s star.
- For each vertex $i \in V'$ we create a vertex u_i with internal opinion 1.
- For each edge $(i, j) \in E'$ we create directed edges $(v_{i,j}, u_i)$ and $(v_{i,j}, u_j)$.
- We create an isolated node w with internal opinion -3 .

We illustrate part of this construction in Figure 4.




Figure 4: A partial illustration of the construction in Proposition 5.4.

Let T be the set of vertices such that adding edges from all the nodes in T to node w minimizes the cost of the Nash equilibrium. Denote by $G + T$ the graph constructed by adding to G edges from all nodes in T to w . Consider some node b with internal opinion of 0. Observe that since nodes with internal opinion of 0 do not have any incoming edges, any edge that would be added from b to w would only effect b 's cost. Since adding an edge from b to w amounts to adding a positive term to b 's cost function, it cannot be the case that this improves b 's cost and thus the total social cost. Thus, T contains only vertices of type $v_{i,j}$ and u_i . In the table in Figure 5 we compute $v_{i,j}$'s opinion in the Nash equilibrium and the cost of its star as a function of which vertices that influence $v_{i,j}$ are in T . For example in the first row we consider the case in which $v_{i,j}, u_i, u_j \notin T$. In this case, $v_{i,j}$'s opinion is $(1 + 1 + 1)/3 = 1$ and the cost of its star is $\frac{1}{2} \cdot 24 = 12$. We use the costs in this table to reason about the structure of T and the cost of the Nash equilibrium in $G + T$.

	Configuration	$v_{i,j}$'s opinion	$v_{i,j}$'s star cost
1	$v_{i,j}, u_i, u_j \notin T$	1	12
2	$v_{i,j} \in T, u_i, u_j \notin T$	0	12
3	$v_{i,j}, u_i \in T, u_j \notin T$	$-1/2$	14
4	$v_{i,j}, u_i, u_j \in T$	-1	20
5	$v_{i,j}, u_j \notin T, u_i \in T$	$1/3$	4
6	$v_{i,j} \notin T, u_i, u_j \in T$	$-1/3$	4

Figure 5: The total cost of $v_{i,j}$'s star for different configurations

In Lemma 5.5 we show how to construct from T a set T' such that $c_{G+T'}(x') = c_{G+T}(x)$ and T' is a *pseudo vertex cover*. We say that a set T' is a pseudo vertex cover if it obeys two properties: (i) it contains only vertices of the type u_i . (ii) the vertices in V' corresponding to the u_i 's in V constitute a vertex cover in G' .

Next, we consider the cost of the Nash equilibrium in the graph $G + S$ where S is a pseudo vertex cover: By the table in Figure 5 we have the cost associated with every $v_{i,j}$'s star is 4. This is by the fact that S is a pseudo vertex cover and hence the only applicable cases are 5 and 6, in both cases the total cost of $v_{i,j}$'s star is 4. Also, note that the cost for each $u_i \in S$ is 8. Hence, the total cost of the Nash equilibrium for network $G + S$ is $f(S) = 4|E| + 8|S|$. By construction, T' is a pseudo vertex cover and it also minimizes $f(\cdot)$, since $c_{G+T}(x) = c_{G+T'}(x') = 4|E| + 8|T'|$ and T is optimal. Therefore T' corresponds to a minimum vertex cover in G' . A key element in this reduction is the property that the cost of $v_{i,j}$'s star is the same, whether $u_i \in T'$ or both u_i and u_j belong to T' .

Lemma 5.5 *There exists a pseudo vertex cover T' such that $c_{G+T'}(x') = c_{G+T}(x)$*

Proof: First, we obtain T'' by removing from T all vertices of type $v_{i,j}$. We have that $c_{G+T''}(x'') \leq c_{G+T}(x)$ since by examining the table in Figure 5 we observe that including vertices of type $v_{i,j}$ in T'' can only increase the cost of the Nash equilibrium. Since T is optimal, it has to be the case that $c_{G+T''}(x'') = c_{G+T}(x)$. Next, to get T' we take T'' and for each vertex $v_{i,j}$ such that $u_i, u_j \notin T''$ we add u_i to T' . By adding these vertices we have not increased the cost since in the worst case $v_{i,j}$'s star and u_i have a total cost of 12 which is the same as their previous total cost. As before by the optimality of T we could not have reduced the cost by adding the vertices, therefore it still holds that $c_{G+T'}(x') = c_{G+T}(x)$. To complete the proof observe that by construction T' is a pseudo vertex cover. □

□

Adding an arbitrary set of edges

In the last case we consider, which is the most general one, we can add any set of edges. For this case we leave open the question of the hardness of adding an unrestricted 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 [14] :

Proposition 5.6 *Finding a best set of arbitrary k edges is NP-hard.*

Proof: We show a reduction from the “Dense k -Subgraph Problem” defined in [14]: given an undirected graph $G' = (V', E')$ and a parameter k , find a set of k vertices with maximum average

degree in the subgraph induced by this set. Given an instance of the “Dense k -Subgraph Problem” we create an instance of the opinion game as follows: (illustrated in Figure 6)

- For every edge $(i, j) \in E'$ we create a node $v_{i,j}$ with internal opinion 0.
- For every vertex $i \in V'$ we create a node u_i with internal opinion 1.
- For every $v_{i,j}$ we add directed edges $(v_{i,j}, u_i)$ and $(v_{i,j}, u_j)$.
- For every u_i we create an in-directed star with $2n^3$ peripheral nodes that have an internal opinion of 0. We later refer to node u_i and all the nodes directed to it as u_i 's star.
- Finally, we create a single isolated vertex w with internal opinion -1 .




Figure 6: A partial illustration of the construction in Proposition 5.6.

The proof is composed of two lemmas. In Lemma 5.7 we show that all edges in the minimizing set are of type (u_i, w) . Then we denote by T the set of nodes of type u_i such that adding an edge from each one of these nodes to w minimizes the cost, and in Lemma 5.8 we show that T is a k densest subgraph.

Lemma 5.7 *The best set of edges to add contains only edges from nodes of type u_i to w .*

Proof: We first show that in the optimal solution the sum of costs exhibited by the nodes of k u_i 's stars should be reduced each by $\Theta(n^3)$. Next, we show that such a cost reduction can only be achieved by connecting k nodes of type u_i to w .

Let F be the set of edges added in some optimal solution. Denote by M the k stars exhibiting the lowest cost in equilibrium after adding the edges in F . Let u_j 's star be the star with the maximal associated cost in M and let a be its cost. Let $b = n^3 - a$ be the cost reduction exhibited by u_j 's star; this is the *minimal* cost reduction exhibited by a star in M . The total cost reduction from adding the edges in F is at most $b + (k - 1)n^3 + n^2$. The reason for this is that any choice of k edges can affect the cost of at most k stars (each has initial cost of n^3) and at most n^2 nodes of type $v_{i,j}$ (each has initial cost of less than 1). On the other hand, the cost reduction of connecting k nodes of type u_i to w is at least $k \cdot (n^3 - 2)$. By comparing the two cost reductions we have that $b > n^3 - n^2 - 2n$. In particular, for, $n \geq 3$ this implies that any optimal solution should reduce the cost of each of k stars by at least 12. The graph structure implies that this can only be done by adding edges to influence the nodes of type u_i to hold opinions closer to the opinion of their

peripheral nodes. Thus, we have that each of the k edges should originate from a distinct node of type u_i .

Next, observe that the opinion u_i holds in equilibrium is $\frac{1}{2}$ + half the opinion its new directed neighbor holds. Therefore, if the opinion of u_i 's new directed neighbor is non-negative then u_i 's opinion is at least $\frac{1}{2}$. In such a case the cost associated with u_i 's star is at least $\frac{1}{8} \cdot 2n^3$. Therefore, the cost reduction is at most $\frac{3}{4} \cdot n^3$. Recall, that the cost reduction for each of k stars should be at least $n^3 - n^2 - 2n$. Thus, for $n \geq 6$ this cost reduction is less than the minimal required cost reduction. The proof is completed by observing that by just adding edges from nodes of type u_i to other nodes it is impossible to create a node (other than w) that holds a negative opinion in equilibrium. □

Lemma 5.8 *The previously defined set T is a solution to the dense k -subgraph problem.*

Proof: The key point is the fact that the cost associated with a node of type $v_{i,j}$ is 0 if and only if both u_i and u_j are in T ; otherwise this cost is exactly $\frac{2}{3}$. When $u_i \in T$, the opinion of u_i in the Nash equilibrium is 0 since it is averaging between 1 and -1 . Therefore node $v_{i,j}$'s associated cost in the Nash equilibrium is:

- 0 - if both u_i and u_j are in T - since $v_{i,j}$ holds opinion 0.
- $\frac{2}{3}$ - if both u_i and u_j are not in T - since $v_{i,j}$'s opinion is $\frac{2}{3}$ and therefore the cost is $(0 - \frac{2}{3})^2 + 2(1 - \frac{2}{3})^2 = \frac{2}{3}$.
- $\frac{2}{3}$ - if only one of u_i, u_j is in T - then $v_{i,j}$'s opinion is $\frac{1}{3}$ and therefore the cost is $(0 - \frac{1}{3})^2 + (0 - \frac{1}{3})^2 + (1 - \frac{1}{3})^2 = \frac{2}{3}$.

Hence to minimize the cost of the Nash equilibrium we should choose a set T maximizing the number of nodes of type $v_{i,j}$ for which both u_i and u_j are in T . In the graph G' from the k -dense subgraph problem that set T is a set of vertices and what we are looking for is the set T with an induced graph that has the maximum number of edges. By definition this set is exactly a k -densest subgraph. □

Finding approximation algorithms for all of the problems discussed in propositions 5.2, 5.4, and 5.6 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 5.9 $c_{G'}(x') \leq \frac{9}{4}c_G(y)$.

Proof: By Theorem 3.1 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 $c_{G'}(x') \leq \frac{9}{4}c_G(y)$. □

Observe that for undirected graphs, we can only include the reverse copies of edges that are not already a part of G . For example if a graph contains both edges (v, u) and (u, v) we cannot add edges to it in order to construct a new graph G' which A is the Laplacian of G' . Once we deal with weighted graphs the restriction no longer holds and therefore by including reverse copies of all edges that do appear in G we can achieve an approximation ratio of 2.

5.1 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 [17, §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 5.10 *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 . As we have seen in Claim 2.2, 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. \square

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 5.11 *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 $\phi \rightarrow \phi_\infty = (x_i - x_j)/(v_{i,i} - v_{i,j})$ 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_∞ (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_∞ , or at β_{ij}/α_{ij} (assuming this point is between 0 and ϕ_∞). We can therefore determine the optimal weight for a single edge in polynomial time. \square

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.

Finding the optimal choice of additional edge weights is a smooth inequality-constrained optimization problem, so while computing a global optimum may be NP-hard, standard gradient projection methods are guaranteed to converge to a constrained stationary point [7, Chapter 2].

Acknowledgments

We thank Michael Macy for valuable discussion.

References

- [1] D. Acemoglu, A. Ozdaglar and A. ParandehGheibi, “Spread of (Mis)Information in Social Networks” in *Games and Economic Behavior*, 2010.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] L. Backstrom, J. Kleinberg, R. Kumar, “Optimizing Web traffic via the media scheduling problem,” *Proc. ACM SIGKDD Conf. Knowledge Disc. Data Mining*, 2009.
- [6] E. Bakshy, I. Rosenn, C. A. Marlow, L. A. Adamic. The Role of Social Networks in Information Diffusion. *Proc. Internationla World Wide Web Conference*, 2012.
- [7] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific 1999.
- [8] K. Bhawalkar, S. Gollapudi, K. Munagala. Coevolutionary opinion formation games. *Proc. ACM Symposium on Theory of Computing*, 2013.
- [9] Y. Boykov, O. Veksler, and R. Zabih, “Fast approximate energy minimization via graph cuts,” in *Proc. 7th Intl. Conf. on Computer Vision*, 1999.
- [10] F. Chierichetti, J. Kleinberg, S. Oren. On discrete preferences and coordination. *Proc. 14th ACM Conference on Electronic Commerce*, 2013.
- [11] F. R. K. Chung, *Spectral Graph Theory*. AMS 1997.
- [12] M. H. DeGroot, “Reaching a consensus,” *J. American Statistical Association*, vol. 69, pp. 118–121, 1974.
- [13] P. M. DeMarzo, D. Vayanos, and J. Zweibel, “Persuasion bias, social influence, and unidimensional opinions,” *Quarterly Journal of Economics*, vol. 118, no. 3, 2003.
- [14] U. Feige, D. Peleg, and G. Kortsarz, “The dense k -subgraph problem,” *Algorithmica*, vol. 29, no. 3, 2001.
- [15] N. E. Friedkin and E. C. Johnsen, “Social influence and opinions,” *J. Math. Sociology*, vol. 15, no. 3-4, 1990.
- [16] B. Golub and M. O. Jackson, “Naive learning in social networks: Convergence, influence and the wisdom of crowds,” *American Econ. J.: Microeconomics*, 2(2010).
- [17] G. H. Golub and C. F. V. Loan, *Matrix Computations*, 3rd ed. Johns Hopkins University Press, 1996.
- [18] P. Groeber, J. Lorenz, F. Schweitzer. Dissonance minimization as a microfoundation of social influence in models of opinion formation. *J. Math. Sociology*, to appear.

- [19] M. O. Jackson, *Social and Economic Networks*. Princeton University Press, 2008.
- [20] 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.
- [21] D. Krackhardt, “A plunge into networks,” *Science*, vol. 326, pp. 47–48, 2 October 2009.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] R. Varga, *Matrix Iterative Analysis*. Prentice Hall, 1962.