HOW BAD IS FORMING YOUR OWN OPINION?

Sigal Oren

Joint work with David Bindel and Jon Kleinberg

How do People Form their Opinion?

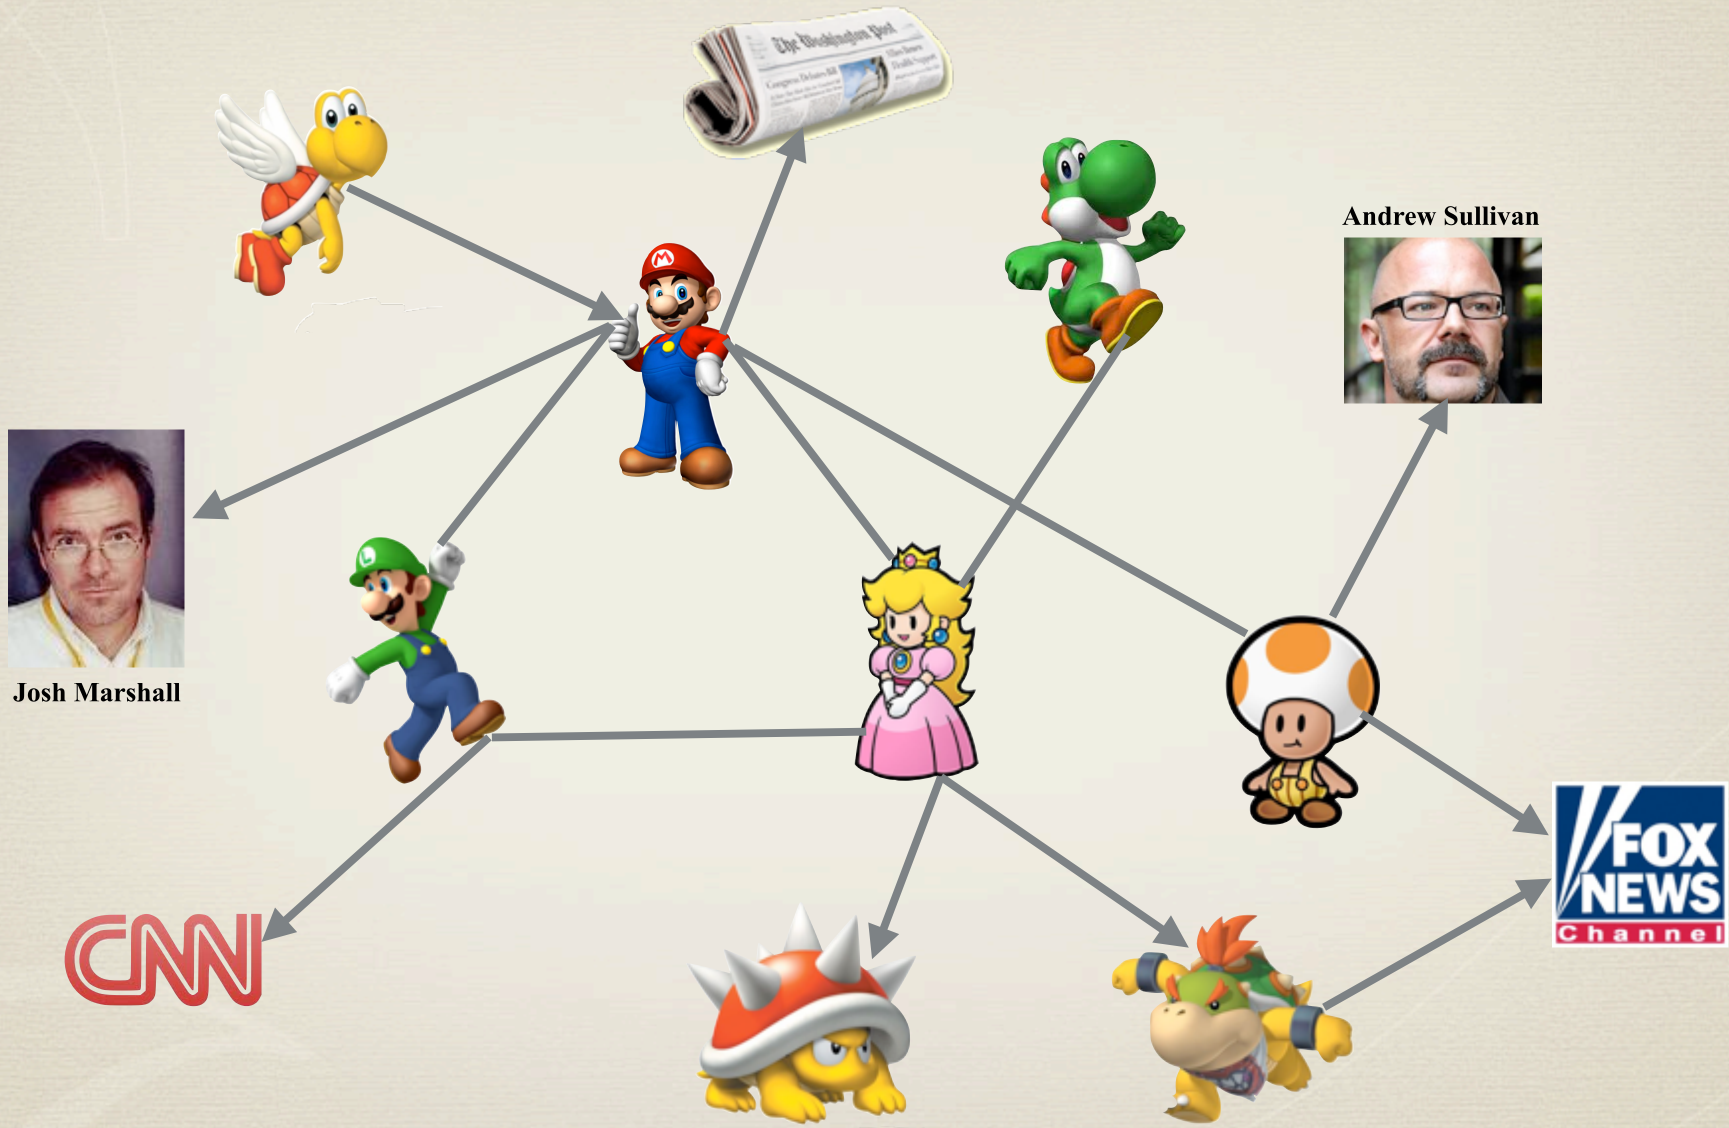

Modeling Opinions

* Opinions are modeled by numeric continuous values.
* One-dimensional :
  * Correlation of political beliefs.
  * Standard model in the literature.
* Examples of opinions that might have continues values:
  * Location on the political spectrum
  * Tax rates
  * Probability of some event happening

DeGroot Model (1974)

Dynamic process of how people form their opinion:

- G=(V,E) - A social network (possibly weighted)
- At step t agent i holds opinion $z_i(t)$ -(initial opinions $z_i(0)$)
- Update rule: At time step t+1 player i updates his opinion to be the weighted average of his and his neighbors opinions at step t.

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

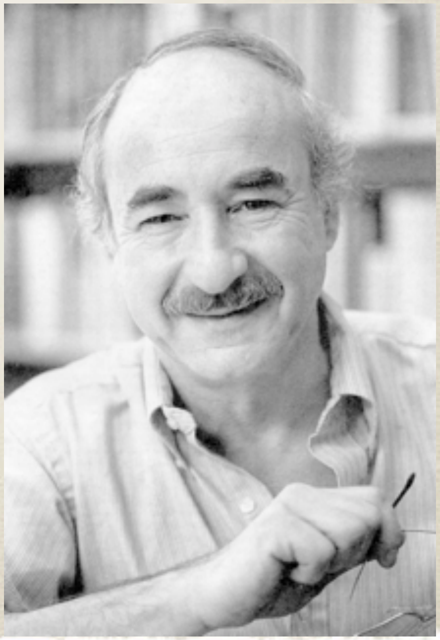

Converges to a consensus under some mild assumptions.

Related work - Analyzes this model and variants, mainly attempts to answer questions such as whether the process converges to a consensus or not [DeMarzo et al (2003), Hegselmann and Krause (2002) ,Lorenz(2005) Golub and Jackson 2007 .]

However...

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"

Friedkin and Johnsen (1990): Each agent has an internal opinion (s_i) which doesn't change.

New update rule: $z_i = \dfrac{s_i + \sum_{j \in N(i)} w_{i,j} z_j}{1 + \sum_{j \in N(i)} w_{i,j}}$

Opinion Formation Game:

- Every player i chooses to hold opinion z_i.
- Player i's cost is $c_i(z) = \underbrace{(z_i - s_i)^2}_{\substack{\text{discomfort from being} \\ \text{far from his internal opinion}}} + \underbrace{\sum_{j \in N(i)} w_{i,j}(z_i - z_j)^2}_{\substack{\text{discomfort from being} \\ \text{far from his friends opinions}}}$

Observation: the Friedkin and Johnsen update rule minimizes the player's cost in this game (best response).

$$c_i'(z) = 2(z_i - s_i) + 2\sum_{j \in N(i)} (z_i - z_j) = 0 \implies z_i = \frac{s_i + \sum_{j \in N(i)} w_{i,j} z_j}{1 + \sum_{j \in N(i)} w_{i,j}}$$

Claim: Repeated averaging converges to the Nash equilibrium which is unique and always exists.

- Social cost: $c(z) = \sum_i [(z_i - s_i)^2 + \sum_{j \in N(i)} w_{i,j}(z_i - z_j)^2]$

Example: undirected graphs

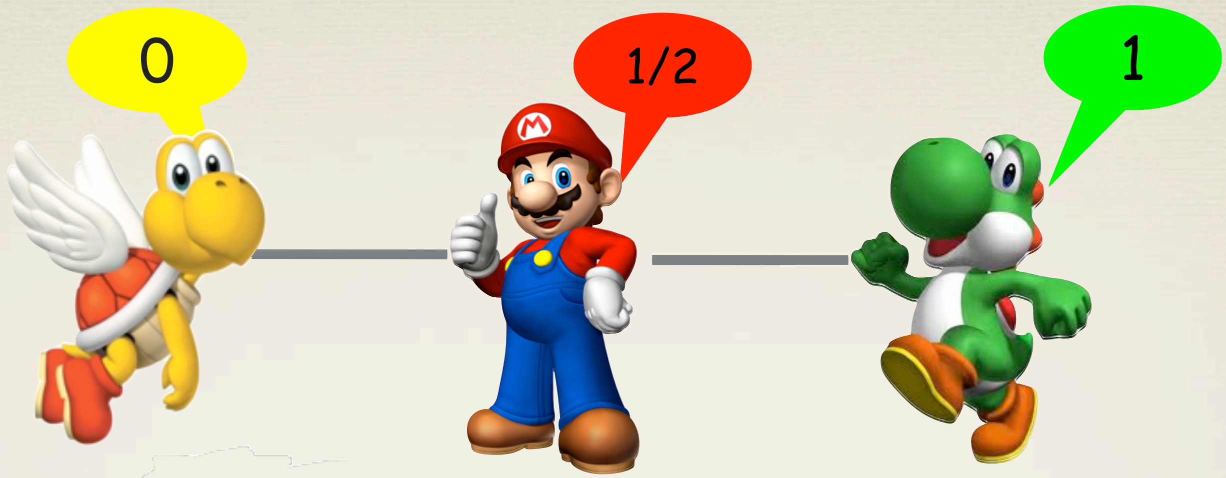

				Total cost:
Nash:	1/4	1/2	3/4	
Cost:	1/8	1/8	1/8	3/8
Optimal:	1/3	1/2	2/3	1/3

$$(\frac{1}{4}-0)^2+(\frac{1}{4}-\frac{1}{2})^2=\frac{1}{8}$$

Theorem: For any undirected graph, the price of anarchy is bounded by 9/8. The bound of 9/8 is tight.

Price of Anarchy is 9/8

Outline

- Undirected Graphs
  - Proof of theorem 1: PoA is bounded by 9/8.
- Directed graphs: can model how people are influenced by media sources.
  - PoA can be unbounded.
  - Families of graphs for which we can get some bounds on the PoA.
- Design: how can we reduce the cost of the Nash equilibrium by modifying the graph?

Computing the Optimal Solution:

Social cost function for undirected graphs:

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

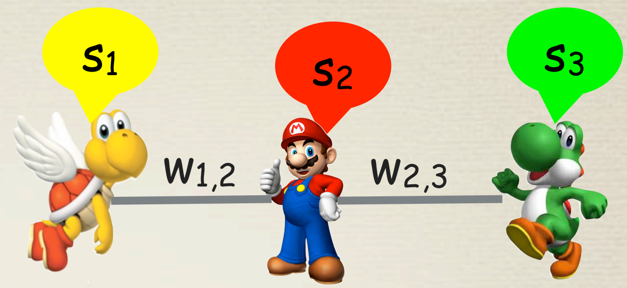

Finding the optimal solution:

$$\frac{\partial c}{\partial z_i} = 2\sum_{j\in N(i)} 2w_{i,j}(z_i - z_j) + 2(z_i - s_i) = 0$$

$$\underbrace{\sum_{j\in N(i)} 2w_{i,j}z_i + \sum_{j\in N(i)} -2w_{i,j}z_j}_{2Lz} + \underbrace{\sum_{j\in N(i)} z_i}_{Iz} = \underbrace{s_i}_{s} \quad \Longrightarrow \quad (2L+I)z = s$$

$$\Downarrow z = (2L+I)^{-1}s$$

$$L = \begin{array}{|c|c|c|c|}\hline \sum_{j\in N(1)} w_{1,j} & -w_{1,2} & \dots & -w_{1,n} \\ \hline -w_{2,1} & \ddots & & \vdots \\ \hline \vdots & & \ddots & \vdots \\ \hline -wn,1 & \dots & \dots & \sum_{j\in N(n)} w_{n,j} \\ \hline \end{array}$$

The optimal solution is unique and always exists since L is a positive semidefinite matrix and thus $(2L+I)$ is a positive definite matrix.

Undirected Graphs - Matrix Notation

Social cost in matrix notation:

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

$$c(z) = \underbrace{z^T \underbrace{2L}_{A} z}_{\text{stress on edges}} + \underbrace{||z - s||^2}_{\text{distance from internal opinion}}$$

$$L = \begin{array}{|c|c|c|c|} \hline \sum_{j\in N(1)} w_{1,j} & -w_{1,2} & \dots & -w_{1,n} \\ \hline -w_{2,1} & \ddots & & \vdots \\ \hline \vdots & & \ddots & \vdots \\ \hline -wn,1 & \dots & \dots & \sum_{j\in N(n)} w_{n,j} \\ \hline \end{array}$$

We have that the optimal solution is: $o = (A + I)^{-1} s$

By writing down the equations for the Nash equilibrium we get that:

$$x = (\frac{1}{2}A + I)^{-1} s$$

1/2 A+I is a positive definite matrix therefore Nash equilibrium is unique and always exists.

Warm up: PoA is Bounded by 2

Observation: Nash equilibrium $x = (\frac{1}{2}A + I)^{-1}s$ minimizes the function:

$$f(z) = z^T(\frac{1}{2}A)z + ||z - s||^2$$

Reminder: Optimal solution minimizes the function:

$$c(z) = z^T Az + ||z - s||^2$$

Putting it all together:

$$PoA = \frac{c(x)}{c(o)} \leq \frac{2f(x)}{c(o)} \leq \frac{2f(o)}{c(o)} \leq \frac{2c(o)}{c(o)} = 2$$

This only holds for undirected graphs!

PoA≤9/8, Step 1: Diagonalizing

The cost of the optimal solution:

$$c(o) = c((A+I)^{-1}s) = \overbrace{((A+I)^{-1}s)^T A((A+I)^{-1}s) + ||((A+I)^{-1}s) - s}^{B}||^2$$

The cost of the Nash equilibrium:

$$c(x) = c((\frac{1}{2}A+I)^{-1}s) \quad = s^T \underbrace{[(\tfrac{1}{2}A+I)^{-1} - I)^2 + (\tfrac{1}{2}A+I)^{-1}A(\tfrac{1}{2}A+I)^{-1}]}_{C} s$$

Our Goal is to bound: $PoA = \frac{c(x)}{c(o)} = \frac{s^T C s}{s^T B s}$

Writing B and C in a "simpler" way can help us achieve this goal!

Observation: B and C are rational functions of A.

Lemma: Matrices A,B and 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 $C = Q\Lambda^C Q^T$

$$\Lambda^M = \begin{array}{|c|c|c|} \hline \lambda_1^M & 0 & 0 \\ \hline 0 & \ddots & 0 \\ \hline 0 & 0 & \lambda_n^M \\ \hline \end{array}$$

Reminder: the cost function is $c(z) = z^T Az + ||z-s||^2$

PoA≤9/8, Step 2: Simplifying

By the diagonalization: $PoA = \frac{c(x)}{c(o)} = \frac{s^T C s}{s^T B s} = \frac{s^T Q \Lambda^C Q^T s}{s^T Q \Lambda^B Q^T s}$

$$s^T Q \Lambda^C Q^T s : \quad \begin{bmatrix} s_1 & \dots & s_n \end{bmatrix} \bullet \begin{bmatrix} v_1 & \dots & v_n \end{bmatrix} \bullet \begin{bmatrix} \lambda^C{}_1 & 0 & 0 \\ 0 & \dots & 0 \\ 0 & 0 & \lambda^C{}_n \end{bmatrix} \bullet \begin{bmatrix} v_1 \\ \dots \\ v_n \end{bmatrix} \bullet \begin{bmatrix} s_1 \\ \dots \\ s_n \end{bmatrix}$$

define $s' = Q^T s$: $\begin{bmatrix} s'_1 & \dots & s'_n \end{bmatrix} \bullet \begin{bmatrix} \lambda^C{}_1 & 0 & 0 \\ 0 & \dots & 0 \\ 0 & 0 & \lambda^C{}_n \end{bmatrix} \bullet \begin{bmatrix} s'_1 \\ \dots \\ s'_n \end{bmatrix}$

By performing the multiplication we get:

$$PoA = \frac{s'^T \Lambda^C s'}{s'^T \Lambda^B s'} = \frac{\sum_{i=1}^n \lambda_i^C {s'_i}^2}{\sum_{i=1}^n \lambda_i^B {s'_i}^2} \leq \max_i \frac{\lambda_i^C {s'_i}^2}{\lambda_i^B {s'_i}^2} = \max_i \frac{\lambda_i^C}{\lambda_i^B}$$

PoA≤9/8, Step 3: Maximizing

Current State: $PoA \leq \max_i \frac{\lambda_i^C}{\lambda_i^B}$

Let λ_i be an eigenvalue of A. Since B and C are rational functions of A we can compute the eigenvalues of B and C as a function of the eigenvalues of A and get: $\lambda_i^B = \frac{\lambda_i}{(\lambda_i + 1)}$ $\qquad \lambda_i^C = \frac{\lambda_i^2 + 4\lambda_i}{(\lambda_i + 2)^2}$

To get an upper bound on the PoA we should maximize: $\phi(\lambda) = \frac{\lambda^2 + 5\lambda + 4}{\lambda^2 + 4\lambda + 4}$

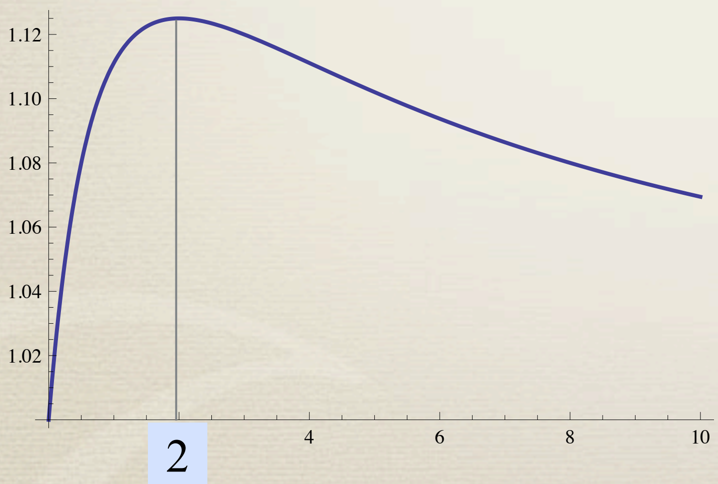

$\Phi(\lambda)$ is maximized for $\lambda=2$ and its value at this point is 9/8. Thus PoA≤9/8.

Given a specific graph, which internal opinions vector maximizes the PoA?

Let λ_j be the eigenvalue of matrix A

maximizing $\phi(\lambda) = \dfrac{\lambda^2 + 5\lambda + 4}{\lambda^2 + 4\lambda + 4}$

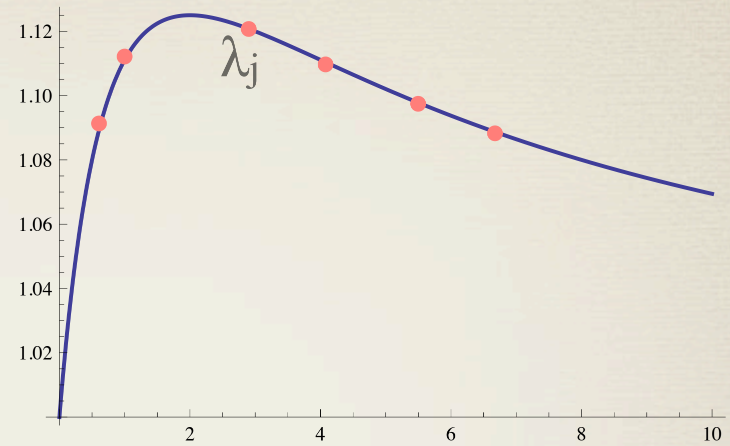

To maximize the PoA, we need for the following bound to hold with equality sign:

$$\frac{\sum_{i=1}^{n} \lambda_i^C {s'_i}^2}{\sum_{i=1}^{n} \lambda_i^B {s'_i}^2} = \max_i \frac{\lambda_i^C}{\lambda_i^B} = \frac{\lambda_j^C}{\lambda_j^B}$$

This can be done by picking $s'_j=1$ and $s'_{i\neq j}=0$.

$$s=Qs': \begin{bmatrix} v_1 & v_j & \dots & v_n \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 1 \\ \dots \\ 0 \end{bmatrix} = v_j$$

Corollary: Graph topology defines an upper bound on PoA. Internal opinions determines the PoA in the predefined range.

Directed Graphs - Star Example

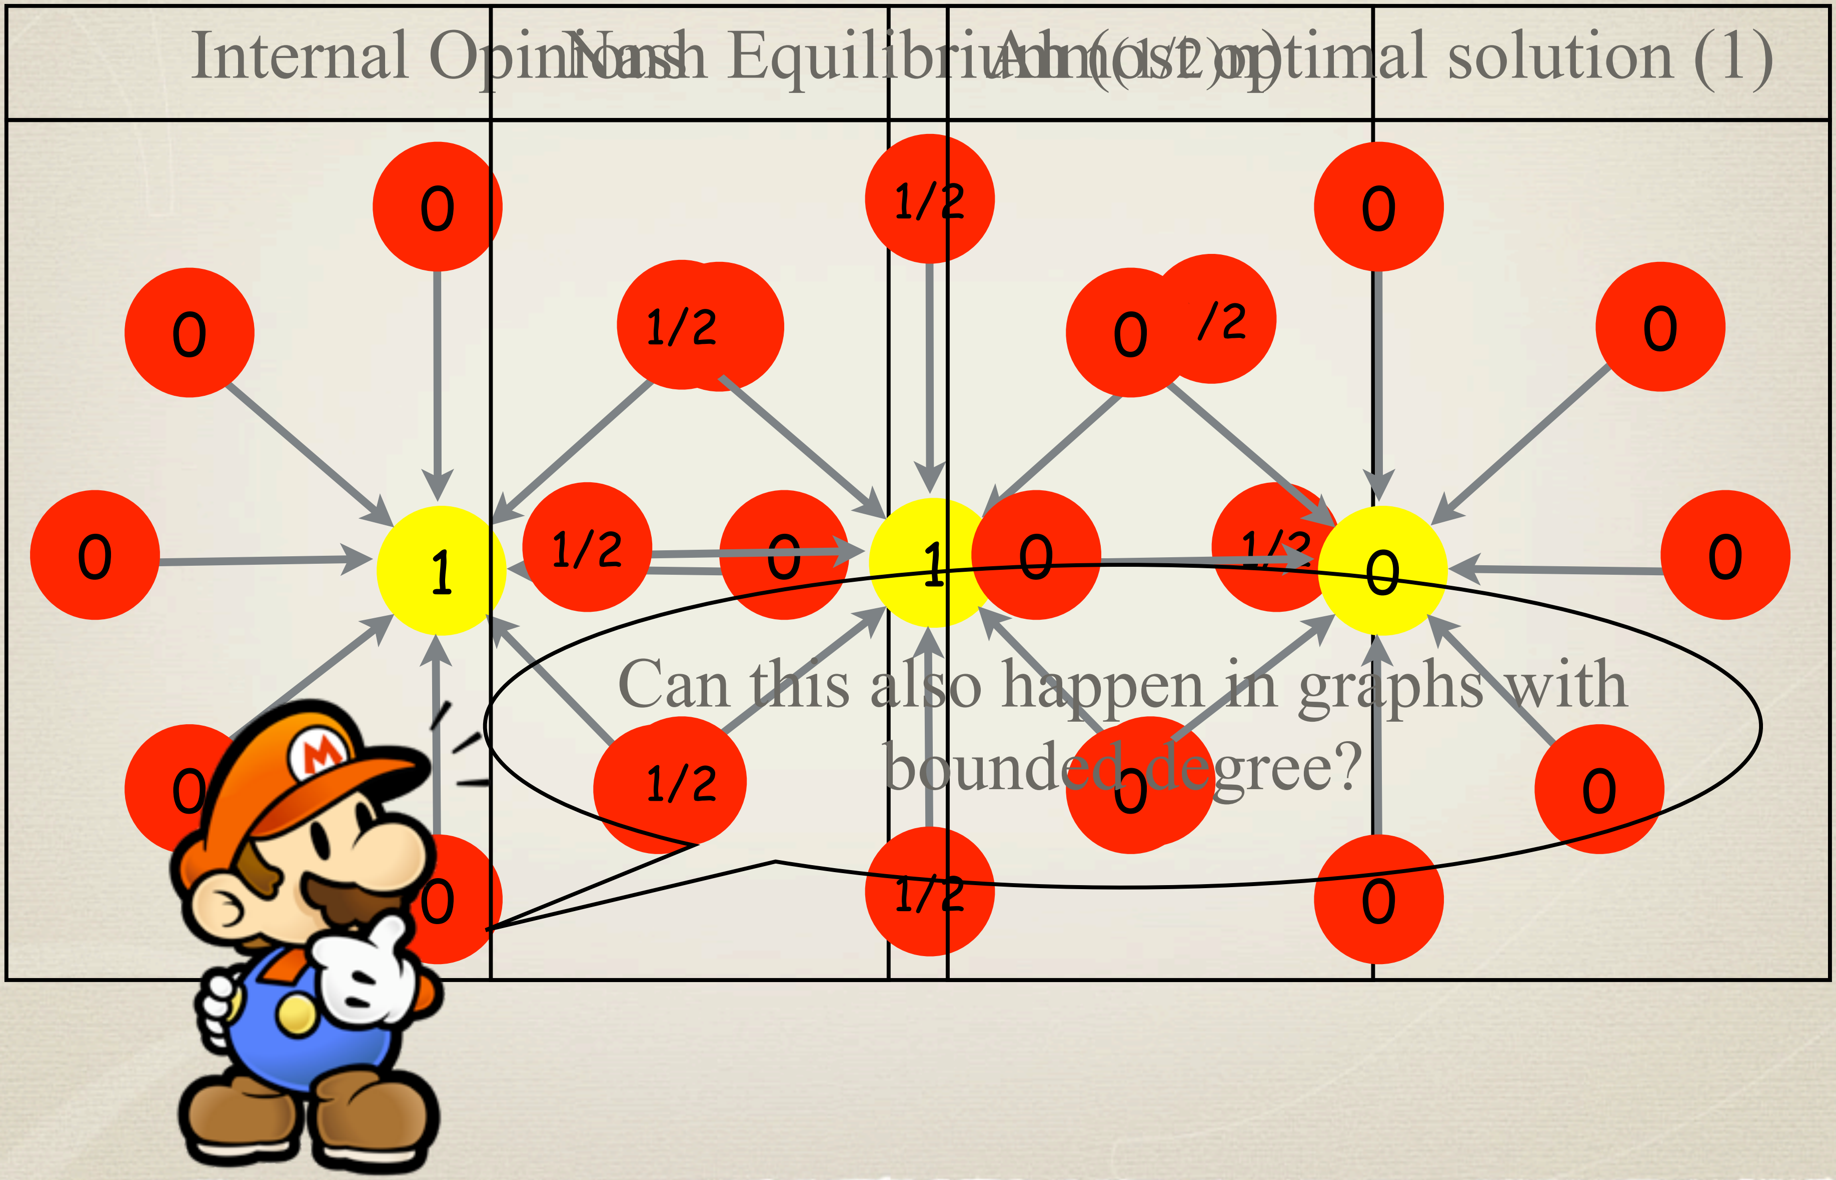

Example: a Tree

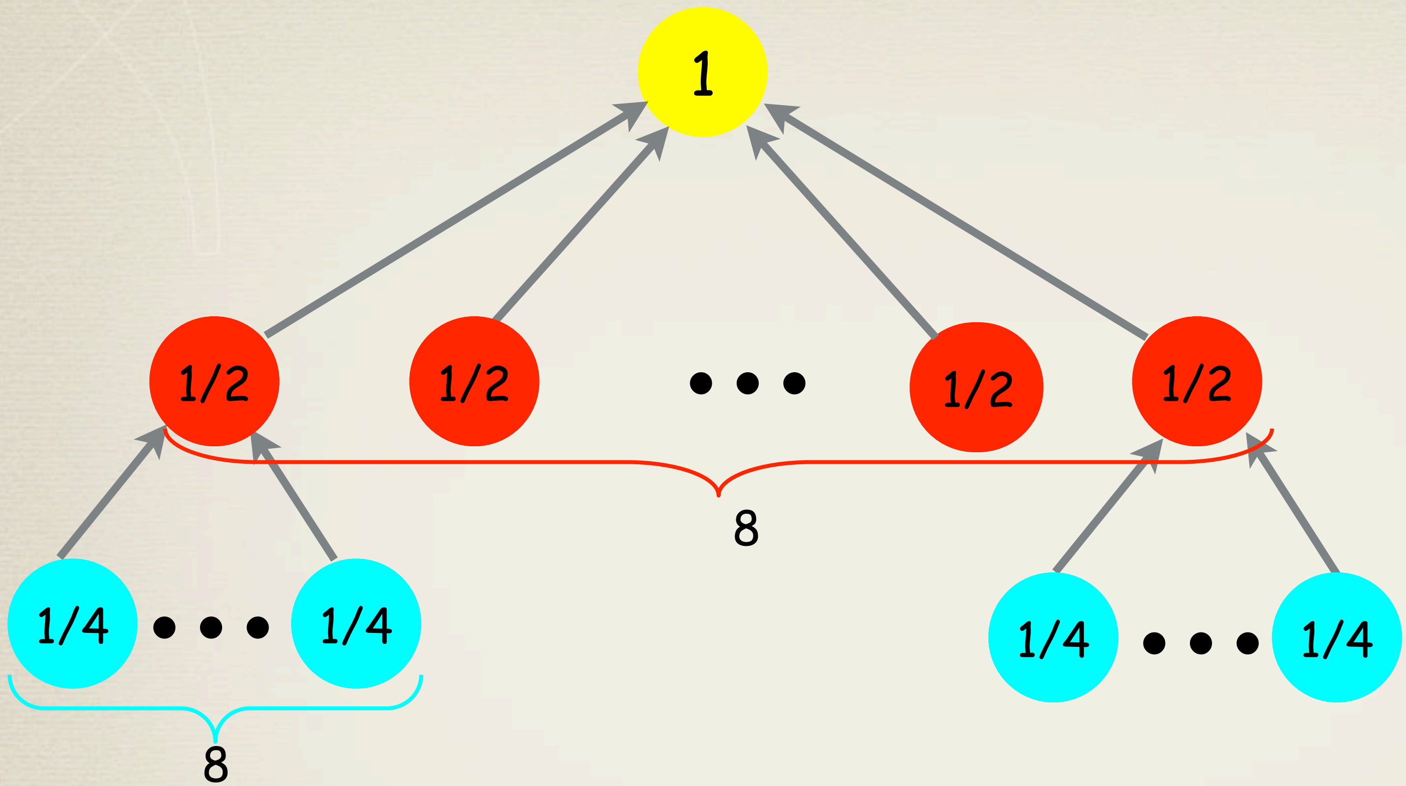

$8 \cdot (2 \cdot (1/2)^2)$

$8^2 \cdot (2 \cdot (1/4)^2)$

In the Nash equilibrium, the total cost of layer i: $2^{3i} \cdot (2 \cdot 2^{-2i}) = 2^{i+1}$

By taking a sum we have that the total cost is: $4(n^{1/3} - 1)$

Optimal solution cost is smaller than 1.

Directed Graphs

Social cost is: $c(z) = z^T A z + ||z - s||^2$

A is now: $$A_{i,i} = \sum_{j \in N(i)} (w_{i,j} + w_{j,i})$$
$$A_{i,j} = -w_{i,j} - w_{j,i}$$

Original Graph

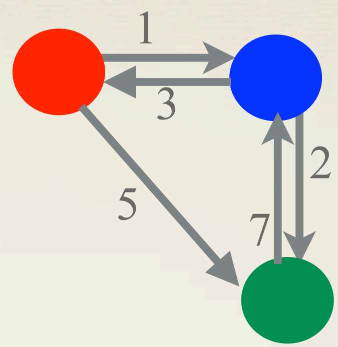

A is the Laplacian of this graph

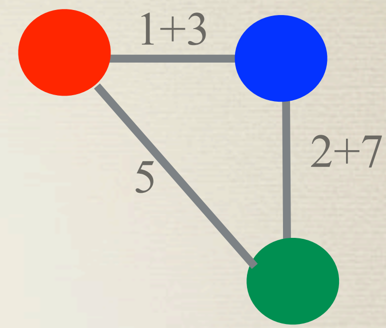

Theorem 2: Given a graph G it is possible to find the maximal PoA and the internal opinions vector yielding it in polynomial time.

- The proof is a generalization of the proof for undirected graphs.
- We can define matrices B and C as before. However, they are no longer a rational function of A.
- Therefore, we can only find the worst internal opinions for a specific graph and not the graph which maximizes the PoA.

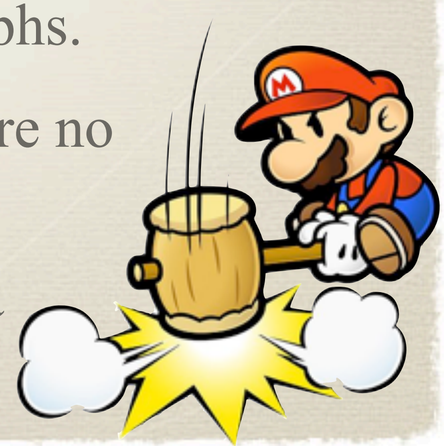

Bounding the PoA: Directed Cycle

Claim 1: For every directed cycle G and any internal opinions vector s: $c(x) \leq \min_z 2(z^T A z) + ||z - s||^2$

Claim 2: The price of anarchy of directed cycles is bounded by 2.

Proof: Let z_1 be the minimizer of $2(z^T A z) + ||z - s||^2$

$$PoA = \frac{c(x)}{c(o)} \leq \frac{2(z_1^T A z_1) + ||z_1 - s||^2}{(o^T A o) + ||o - s||^2}$$

$$\leq \frac{2\lambda_2 + 2}{2\lambda_2 + 1} \leq 2$$

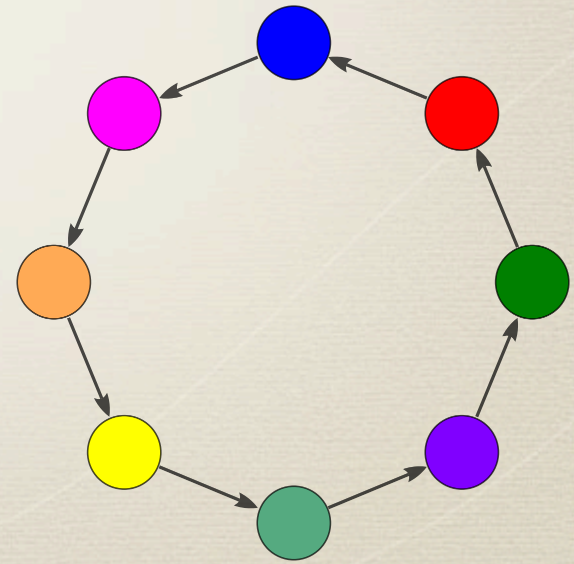

* where λ_2 is the second smallest eigenvalue.

Bounding the PoA: Generalizing

Claim: Let $\mathcal{G}$ be a graph family for which there exists a β such that for any $G \in \mathcal{G}$ and any internal opinions vector s: $c(x) \leq \min_z(\beta z^T Az + ||z-s||^2)$

then, for all $G \in \mathcal{G}$ and s: $PoA \leq \dfrac{\beta\lambda_2 + \beta}{\beta\lambda_2 + 1}$.

(intuition) When does such a β exists?

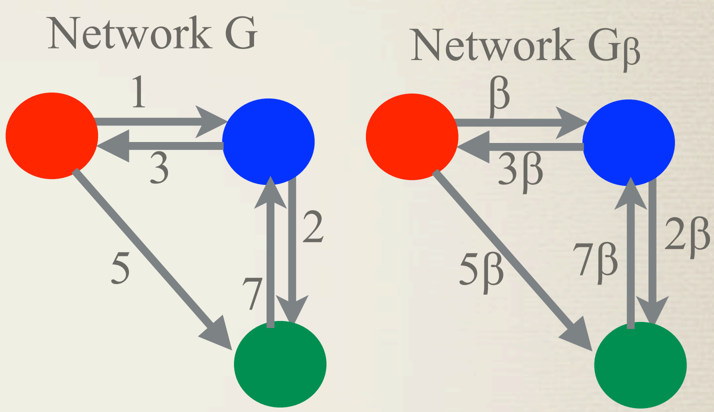

Let z_1 be the minimizer of $\beta z^T Az + ||z-s||^2$, then z_1 is the optimal opinions vector for a different network G_β. Therefore:

$$min_z \beta z^T Az + ||z-s||^2 \leq \underbrace{min_{\{z|\forall i z_i = z_{i+1}\}} ||z-s||^2}_{\boxed{\text{cost of best consensus}}}$$

Corollary: β exists if and only if the cost of the Nash equilibrium is smaller than the cost of the best consensus.

Bounding the PoA: Continue

* But we still don't know how to find the value of the β.
* To get the value of β we use an intermediate function g(z) such that $c(x)=\min_z g(z)$. This function has a specific structure which makes it easier get a bound:

$$g(z) \leq \beta z^T A z + ||z-s||^2$$

Using this approach we can prove the following claim:

Claim: For bounded degree Eulerian graphs:

$$c(x) \leq \min_z((d+1)z^T A z + ||z-s||^2)$$

Therefore, the price of anarchy of Eulerian graphs is bounded by their maximal degree. tight?

Eulerian Graphs and Expansion

More specifically, for Eulerian graphs we have that: $PoA \leq \frac{\beta\lambda_2 + \beta}{\beta\lambda_2 + 1} \leq \beta$

Can we get a bound that doesn't depend on β?

Claim: For an Eulerian asymmetric expander with bounded degree d and expansion α, $PoA \leq O(d^2/\alpha^2)$.

Asymmetric - for every two nodes i,j in the graph, at most one of the two edges (i,j) and (j,i) is part of the graph.

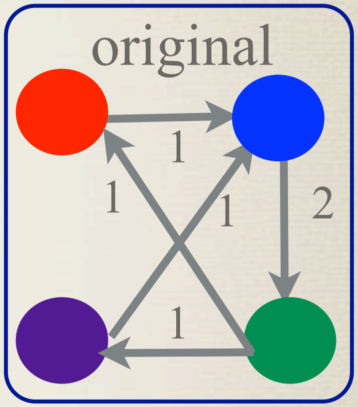

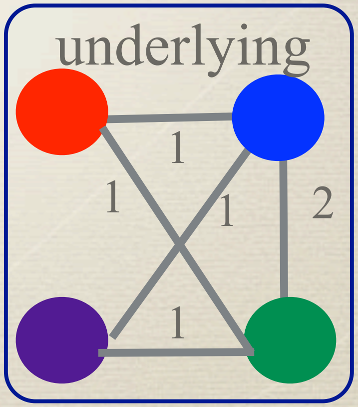

Proof: $$PoA \leq \frac{\beta + \beta\lambda_2}{1 + \beta\lambda_2} \leq \frac{\beta + \beta\lambda_2}{\beta\lambda_2} \leq \frac{1 + \lambda_2}{\lambda_2}$$

$$\leq \frac{2d(1 + 2d)}{\alpha^2} = O(d^2/\alpha^2)$$

From spectral theory we have that:

1. $\lambda_2 \leq \lambda_n \leq 2d$
2. $\lambda_2 \geq \alpha^2/2d$

Improving Equilibrium's Cost

Goal: Given an unweighted graph G, add edges to create a new graph G' such that the cost of the Nash equilibrium is minimal.

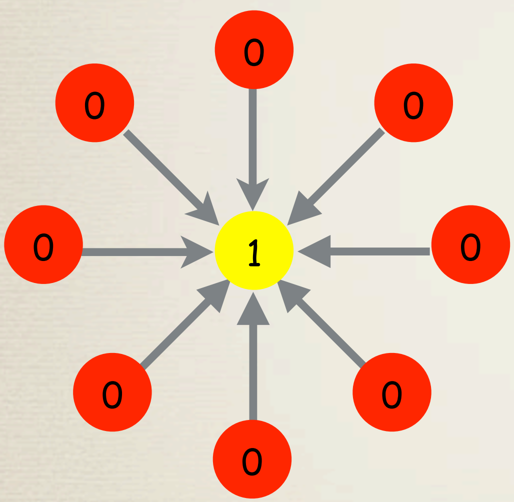

(1/2)n

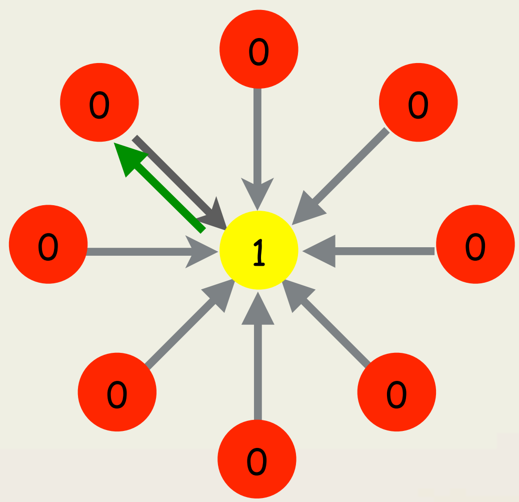

(2/9)n

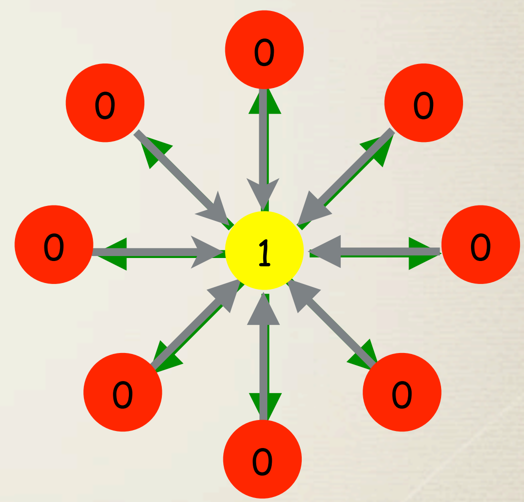

<1

Doubling the Edges

Claim: By making the graph symmetric we can get a 9/4-approximation to the original optimal solution.

Proof: In the beginning of the talk we proved $c_{G'}(x') \leq \frac{9}{8} c_{G'}(o')$
In the worst case we double all edges therefore $c_{G'}(o') \leq 2 c_G(o)$

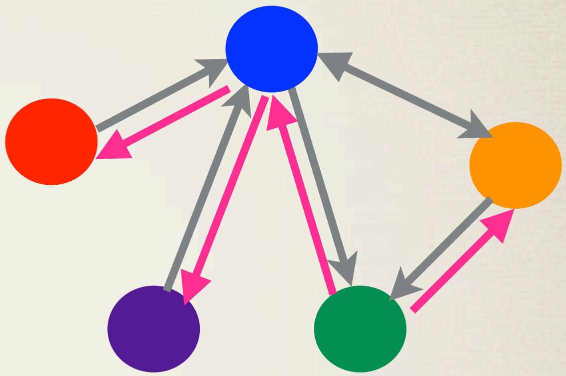

Observation: The cost of the new Nash equilibrium is at least the cost of the optimal solution.

Open question: Can we do better?

Adding Edges by Various Restrictions

1. Best set of edges from a specific node w.
2. Best set of edges to add to a specific node w.
3. Best k edges to add to the graph.

We show that finding the optimal set of edges in all three variants is NP-hard.
Open question: find approximations.

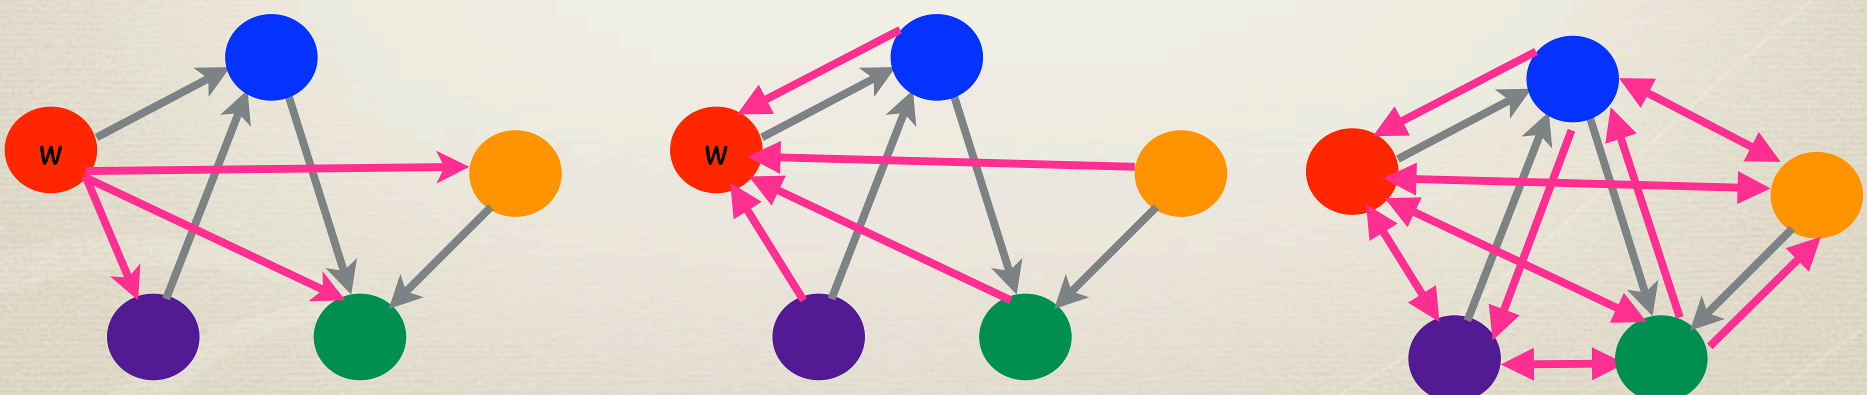

1 2 3

Open Questions

- Can the problem of finding the best set of edges to add (without restrictions) be solved optimally?
- Find approximation algorithms for the NP-hard edge addition problems.
- For bounded degree Eulerian graphs, is the bound of $(d+1)$ tight?
  - Currently we don't have any instance with $PoA>2$

Conclusions

- People are unaware of the externalities of their actions and thus the Nash equilibrium is suboptimal. If the graph is undirected things are not so bad.
- Problems start when the graph is directed -- someone who is very influential doesn't care.
    - For some topologies (specific Eulerian graphs) PoA is still bounded. Is it possible to have a complete characterization?
- It is possible to reduce the cost of the Nash equilibrium by getting people to talk to more people (specific ones).