Controlling Opinion Diffusion on Social Networks

Diodato Ferraioli

Università di Salerno

Progetti di Ricerca 2019

based on joint works with

Vincenzo Auletta, Iannis Caragiannis, Valeria Fionda, Clemente Galdi, Gianluigi Greco, and Giuseppe Persiano

Opinion vs Belief

- Fix a social network service $\boldsymbol{S}$
  - e.g., Facebook, Twitter, Instagram, Google+…

Opinion vs Belief

- Fix a social network service $\boldsymbol{S}$
  - e.g., Facebook, Twitter, Instagram, Google+...
- You believe $\boldsymbol{S}$ is completely messed up

Opinion vs Belief

- Fix a social network service ***S***
  - e.g., Facebook, Twitter, Instagram, Google+. . .
- You believe ***S*** is completely messed up
- You end up in a group of very enthusiast ***S***'s users

Opinion vs Belief

- Fix a social network service ***S***
  - e.g., Facebook, Twitter, Instagram, Google+...
- You believe ***S*** is completely messed up
- You end up in a group of very enthusiast ***S***'s users
- They ask you for your opinion about ***S***

Opinion vs Belief

- Fix a social network service $\boldsymbol{S}$
  - e.g., Facebook, Twitter, Instagram, Google+...
- You believe $\boldsymbol{S}$ is completely messed up
- You end up in a group of very enthusiast $\boldsymbol{S}$'s users
- They ask you for your opinion about $\boldsymbol{S}$
  - opinion $\stackrel{?}{=}$ belief

Private Belief vs Public Opinion in a Social Network

Main theme in Social Sciences:

- an individual has a *private belief*
- the *public opinion* is the result of a strategic decision that depends on the public opinions of the social neighborhood
- the private belief remains constant even as the public opinion is updated

De Groot, J. American Statistical Association, 1974
Friedkin and Johnsen, J. Mathematical Sociology, 1990.

Discrete Preference Games

F, Goldberg, Ventre [SAGT 2012, TCS 2016]
Chierichietti, Kleinberg, Oren [EC 2013]

- Players are vertices of a **social** graph

Discrete Preference Games

F, Goldberg, Ventre [SAGT 2012, TCS 2016]
Chierichietti, Kleinberg, Oren [EC 2013]

- Players are vertices of a **social** graph
- Public **opinions** and private **beliefs** are values in $\{0, 1\}$

Discrete Preference Games

F, Goldberg, Ventre [SAGT 2012, TCS 2016]
Chierichietti, Kleinberg, Oren [EC 2013]

- Players are vertices of a **social** graph
- Public **opinions** and private **beliefs** are values in $\{0, 1\}$
- Cost of player i is

$$C_i(b_i, \mathrm{x}) = \alpha_i \cdot \mid x_i - b_i \mid + (1 - \alpha_i) \cdot \sum_{j \in N(i)} w_{ij} \mid x_i - x_j \mid$$

- $\alpha_i \in [0, 1]$: stubbornness level
- w_{ij}: edge weights

Discrete Preference Games

F, Goldberg, Ventre [SAGT 2012, TCS 2016]
Chierichietti, Kleinberg, Oren [EC 2013]

- Players are vertices of a **social** graph
- Public **opinions** and private **beliefs** are values in $\{0, 1\}$
- Cost of player i is

$$C_i(b_i, \mathrm{x}) = \alpha_i \cdot \mid x_i - b_i \mid + (1 - \alpha_i) \cdot \sum_{j \in N(i)} w_{ij} \mid x_i - x_j \mid$$

- $\alpha_i \in [0, 1]$: stubbornness level
- w_{ij}: edge weights

Properties

- Belief of a player is **given**
- Opinion of a player is outcome of **strategizing**

Our questions

Do opinions converge to some stable state?

Our questions

Do opinions converge to some stable state?

How "good" is this stable state?

Convergence to the equilibrium

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)

Convergence to the equilibrium

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)

- Discrete Preference Games are **Potential Games**

Convergence to the equilibrium

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)

- Discrete Preference Games are **Potential Games**
- Nash equilibria can be computed in **polynomial time**

Convergence to the equilibrium

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)

- Discrete Preference Games are **Potential Games**
- Nash equilibria can be computed in **polynomial time**
- Best-Response Dynamics converge in **polynomial time** to a Nash equilibrium in unweighted graphs

Convergence to the equilibrium

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)

- Discrete Preference Games are **Potential Games**
- Nash equilibria can be computed in **polynomial time**
- Best-Response Dynamics converge in **polynomial time** to a Nash equilibrium in unweighted graphs
  - It takes only **pseudo-polynomial** time in weighted graphs

Quality of the equilibrium

Quality of the equilibrium

Objective functions

- Social cost SC = sum of players' costs

Quality of the equilibrium

Objective functions

- Social cost SC = sum of players' costs
- Social welfare SW = sum of players' utility
  - **Utility** of player i is

$$U_i(b_i, \mathrm{x}) = \alpha_i \,(1- \mid x_i - b_i \mid) + (1-\alpha_i) \sum_{j \in N(i)} (1- \mid x_i - x_j \mid)$$

Quality of the equilibrium

Objective functions

- Social cost SC = sum of players' costs
- Social welfare SW = sum of players' utility
  - **Utility** of player i is

$$U_i(b_i, \mathrm{x}) = \alpha_i\,(1- \mid x_i - b_i \mid) + (1-\alpha_i) \sum_{j \in N(i)} (1- \mid x_i - x_j \mid)$$

Quality Measure

- **Price of Anarchy** = SC(worst equilibrium) / SC(optimum)

Price of Anarchy

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)
(Chierichietti, Kleinberg, Oren, EC 2013)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)

Price of Anarchy

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)
(Chierichietti, Kleinberg, Oren, EC 2013)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)

Cost-minimizing Discrete Preference Games

Price of Anarchy is **unbounded**

- All agents have belief $\mathbf{0}$
- All agents declare opinion $\mathbf{1}$

Price of Anarchy

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)
(Chierichietti, Kleinberg, Oren, EC 2013)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)

Cost-minimizing Discrete Preference Games

Price of Anarchy is **unbounded**

- All agents have belief $\mathbf{0}$
- All agents declare opinion $\mathbf{1}$

Utility-maximizing Discrete Preference Games

- Price of Anarchy is at most **2**

Price of Anarchy

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)
(Chierichietti, Kleinberg, Oren, EC 2013)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)

Cost-minimizing Discrete Preference Games

Price of Anarchy is **unbounded**

- All agents have belief **0**
- All agents declare opinion **1**

Utility-maximizing Discrete Preference Games

- Price of Anarchy is at most **2**
  - $U_i \geq \frac{1}{2}(\alpha + (1-\alpha)|N_i|)$ in equilibrium

Price of Anarchy

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)
(Chierichietti, Kleinberg, Oren, EC 2013)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)

Cost-minimizing Discrete Preference Games

Price of Anarchy is **unbounded**

- All agents have belief **0**
- All agents declare opinion **1**

Utility-maximizing Discrete Preference Games

- Price of Anarchy is at most **2**
  - $U_i \geq \frac{1}{2}(\alpha + (1-\alpha)|N_i|)$ in equilibrium
  - $U_i \leq \alpha + (1-\alpha)|N_i|$

Price of Anarchy

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)
(Chierichietti, Kleinberg, Oren, EC 2013)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)

Cost-minimizing Discrete Preference Games

Price of Anarchy is **unbounded**

- All agents have belief $\mathbf{0}$
- All agents declare opinion $\mathbf{1}$

Utility-maximizing Discrete Preference Games

- Price of Anarchy is at most **2**
  - $U_i \geq \frac{1}{2}(\alpha + (1-\alpha)|N_i|)$ in equilibrium
  - $U_i \leq \alpha + (1-\alpha)|N_i|$
- The bound is tight

Price of Anarchy

(F, Goldberg, Ventre, SAGT 2012, TCS 2016)
(Chierichietti, Kleinberg, Oren, EC 2013)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)

Cost-minimizing Discrete Preference Games

Price of Anarchy is **unbounded**

- All agents have belief **0**
- All agents declare opinion **1**

Utility-maximizing Discrete Preference Games

- Price of Anarchy is at most **2**
  - $U_i \geq \frac{1}{2}(\alpha + (1-\alpha)|N_i|)$ in equilibrium
  - $U_i \leq \alpha + (1-\alpha)|N_i|$
- The bound is tight

The equilibria are the same in both versions

A different measure of quality

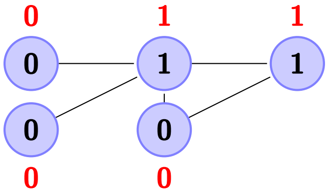

A different measure of quality

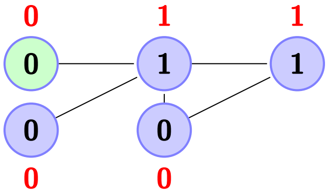

A different measure of quality

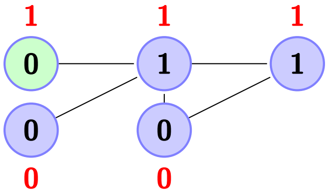

A different measure of quality

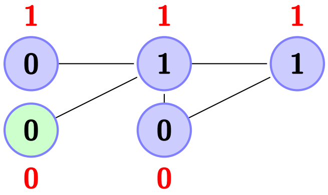

A different measure of quality

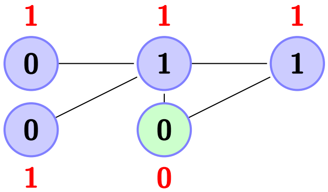

A different measure of quality

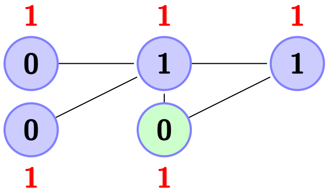

A different measure of quality

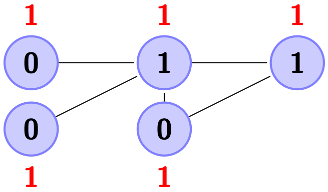

- The distribution of opinions at stable state...
- is **completely different** from the distribution at initial state

Our goal

Our goal

How large can be this difference?

Our goal

How large can be this difference?

How common is it to see a large difference?

Our framework

Worst-case approach

For each graph:

- an adversarially chosen starting state
- an adversarially chosen sequence of updates

Our framework

Worst-case approach

For each graph:

- an adversarially chosen starting state
- an adversarially chosen sequence of updates

Can the adversary subvert the distribution of opinions?

Our framework

Worst-case approach

For each graph:

- an adversarially chosen starting state
- an adversarially chosen sequence of updates

Can the adversary subvert the distribution of opinions?

An useful result [Bredereck and Elkind, 2017]

The number of opinions $\mathbf{1}$ is maximized by **greedy sequences**:

- First flip non-stable agents from $\mathbf{0}$ to $\mathbf{1}$, if any
- Then flip non-stable agents from $\mathbf{1}$ to $\mathbf{0}$, if any

Worst Case Approach?

Is the adversary really able to set both

- starting state (a.k.a., seeds) **and**

Worst Case Approach?

Is the adversary really able to set both

- starting state (a.k.a., seeds) **and**
- the sequence of updates?

Worst Case Approach?

Is the adversary really able to set both

- starting state (a.k.a., seeds) **and**
- the sequence of updates?

Social Proof Marketing

"Hey, join our community.
So many friend of yours are already using our services"

Our results

- Minority can become majority
- A bare majority can become a consensus

For which graphs does minority become majority?

(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2015)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)
(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2017)

For which graphs does minority become majority?

(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2015)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)
(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2017)

The result

Minority can become majority for **all graphs**, except

- very dense graphs ("almost-cliques")
- very sparse graphs ("almost-empty")

For which graphs does minority become majority?

(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2015)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)
(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2017)

The result

Minority can become majority for **all graphs**, except

- very dense graphs ("almost-cliques")
- very sparse graphs ("almost-empty")
- Density thresholds depend on α_i

For which graphs does minority become majority?

(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2015)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)
(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2017)

The result

Minority can become majority for **all graphs**, except

- very dense graphs ("almost-cliques")
- very sparse graphs ("almost-empty")
- Density thresholds depend on α_i

Minus: In our proof, the initial minority is chosen to be large

For which graphs does minority become majority?

(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2015)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)
(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2017)

The result

Minority can become majority for **all graphs**, except

- very dense graphs (“almost-cliques”)
- very sparse graphs (“almost-empty”)
- Density thresholds depend on α_i

Minus: In our proof, the initial minority is chosen to be large

Plus: We prove that the initial minority can be computed in **polynomial time**

For which graphs does minority become majority?

(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2015)
(Auletta, Caragiannis, F, Galdi, Persiano, AAMAS 2017)
(Auletta, Caragiannis, F, Galdi, Persiano, WINE 2017)

The result

Minority can become majority for **all graphs**, except

- very dense graphs ("almost-cliques")
- very sparse graphs ("almost-empty")
- Density thresholds depend on α_i

Minus: In our proof, the initial minority is chosen to be large

Plus: We prove that the initial minority can be computed in **polynomial time**

- Finding a smaller minority that becomes majority is **NP-hard**

From Bare Majority to Consensus

(Auletta, F, Greco, IJCAI 2018)

For **every** graph there is a configuration with $\lceil n/2 \rceil$ **nodes** with opinion $\mathbf{1}$ from which the dynamics converges to **consensus**

From Bare Majority to Consensus

(Auletta, F, Greco, IJCAI 2018)

For **every** graph there is a configuration with $\lceil n/2 \rceil$ **nodes** with opinion $\mathbf{1}$ from which the dynamics converges to **consensus**

The starting configuration can be computed in **poly time**

From Majority to Consensus

Proof Sketch

Partitions and utilities

- The social graph is partitioned in $(\boldsymbol{A}, \boldsymbol{B})$
- Given $(\boldsymbol{A}, \boldsymbol{B})$, the **utility** of a node is the difference between the number of neighbors on her side and on the other side

From Majority to Consensus

Proof Sketch

Partitions and utilities

- The social graph is partitioned in (A, B)
- Given (A, B), the **utility** of a node is the difference between the number of neighbors on her side and on the other side

Nice Partitions

They are partitions from which it is easy to converge to consensus

From Majority to Consensus

Proof Sketch

Partitions and utilities

- The social graph is partitioned in $(\boldsymbol{A}, \boldsymbol{B})$
- Given $(\boldsymbol{A}, \boldsymbol{B})$, the **utility** of a node is the difference between the number of neighbors on her side and on the other side

Nice Partitions

They are partitions from which it is easy to converge to consensus

- They are minimal
- One of the following holds:
  - a node has positive utility
  - a side has all zero-utility nodes with a negative neighbor

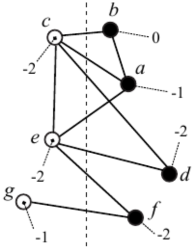

From Majority to Consensus

Proof Sketch – 2

Critical Pair for partition $(X, \bar{X})$

Two nodes x, y on different sides of partition such that, by swapping them, one of the following happens:

- Edges among sides increases
- Nodes that have no negative neighbors on her side decreases
- Distance between these nodes decreases

From Majority to Consensus

Proof Sketch – 2

Critical Pair for partition $(X, \bar{X})$

Two nodes x, y on different sides of partition such that, by swapping them, one of the following happens:

- Edges among sides increases
- Nodes that have no negative neighbors on her side decreases
- Distance between these nodes decreases

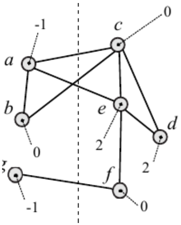

(b, e)

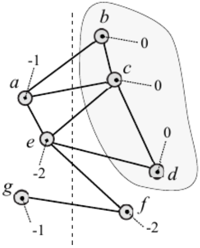

(a, c)

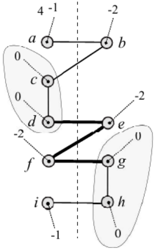

(d, e)

From Majority to Consensus

Proof Sketch – 3

Theorem

If a partition has not a critical pair, then it is nice.

A polynomial algorithm for computing nice partitions

- Start from an arbitrary bisection

From Majority to Consensus

Proof Sketch – 3

Theorem

If a partition has not a critical pair, then it is nice.

A polynomial algorithm for computing nice partitions

- Start from an arbitrary bisection
- **Swap** critical pairs until one exists

From Majority to Consensus

Proof Sketch – 3

Theorem

If a partition has not a critical pair, then it is nice.

A polynomial algorithm for computing nice partitions

- Start from an arbitrary bisection
- **Swap** critical pairs until one exists
  - The number of critical pairs is polynomial in n

Hardness Results

From Minority to Consensus

It is **NP-hard** to decide if a minority converges to consensus

- Reduction from Vertex Cover

Hardness Results

From Minority to Consensus

It is **NP-hard** to decide if a minority converges to consensus

- Reduction from Vertex Cover
- There is a sharp phase transition at $n/2$

What happens if there are **at least three** opinions?

(Auletta, F, Fionda, Greco, AAMAS 2019)

What happens if there are **at least three** opinions?

(Auletta, F, Fionda, Greco, AAMAS 2019)

Greedy sequences fail to maximize opinion $\mathbf{1}$

- In order to maximize $\mathbf{1}$, one needs to interleave the diffusion of different opinions
- It may be necessary to change the opinion of a node $\boldsymbol{n}$ times
  - With two opinions two changes were sufficient

What happens if there are **at least three** opinions?

(Auletta, F, Fionda, Greco, AAMAS 2019)

Greedy sequences fail to maximize opinion **1**

- In order to maximize **1**, one needs to interleave the diffusion of different opinions
- It may be necessary to change the opinion of a node n times
  - With two opinions two changes were sufficient

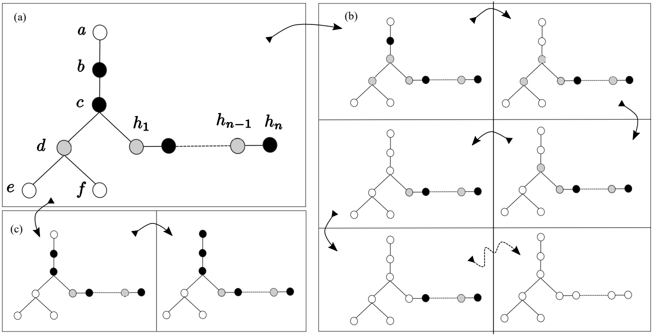

What happens if there are **at least three** opinions?

Majority and Consensus problems are hard

- Problems are still hard if we constrain each agent to update her opinion a limited number of times
- Reduction from One-In-Three Positive 3-SAT

Thank you!