Challenges in Fairness, Accountability, Transparency in AI for Social Good

Fairness in AI for Social Good

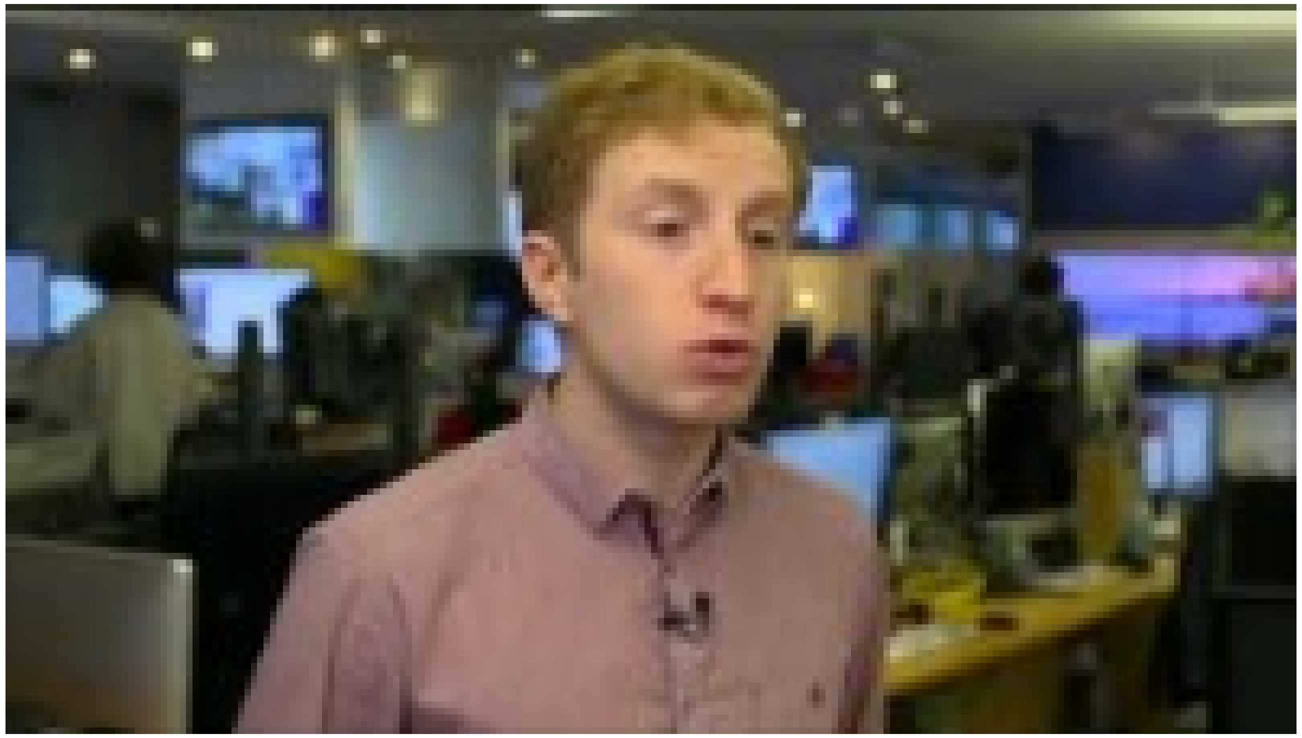

https://www.youtube.com/watch?v=QvRZuHQBTps

What a machine learning tool that turns Obama white can (and can't) tell us about AI bias

A striking image that only hints at a much bigger problem

By James Vincent | Jun 23, 2020, 3:45pm EDT

SHARE

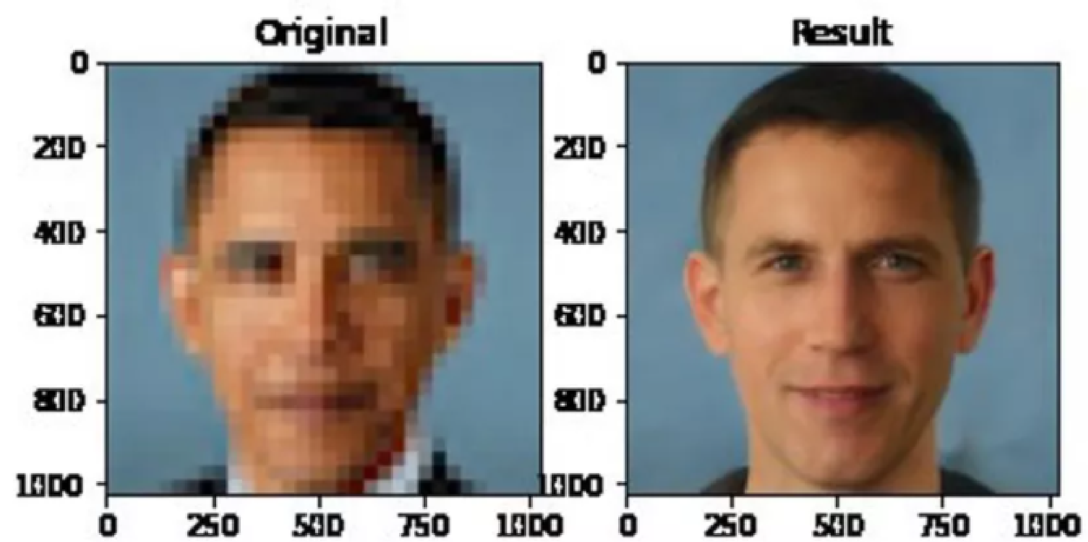

verge deals

Subscribe to get the best Verge-approved tech deals of the week

What PULSE does is use StyleGAN to "imagine" the high-res version of pixelated inputs. It does this not by "enhancing" the original low-res image, but by generating a completely new high-res face that, when pixelated, looks the same as the one inputted by the user.

"It does appear that PULSE is producing white faces much more frequently than faces of people of color," wrote the algorithm's creators on Github. "This bias is likely inherited from the dataset StyleGAN was trained on [...] though there could be other factors that we are unaware of."

The VERGE

Fairness in Predictive vs Prescriptive Phase

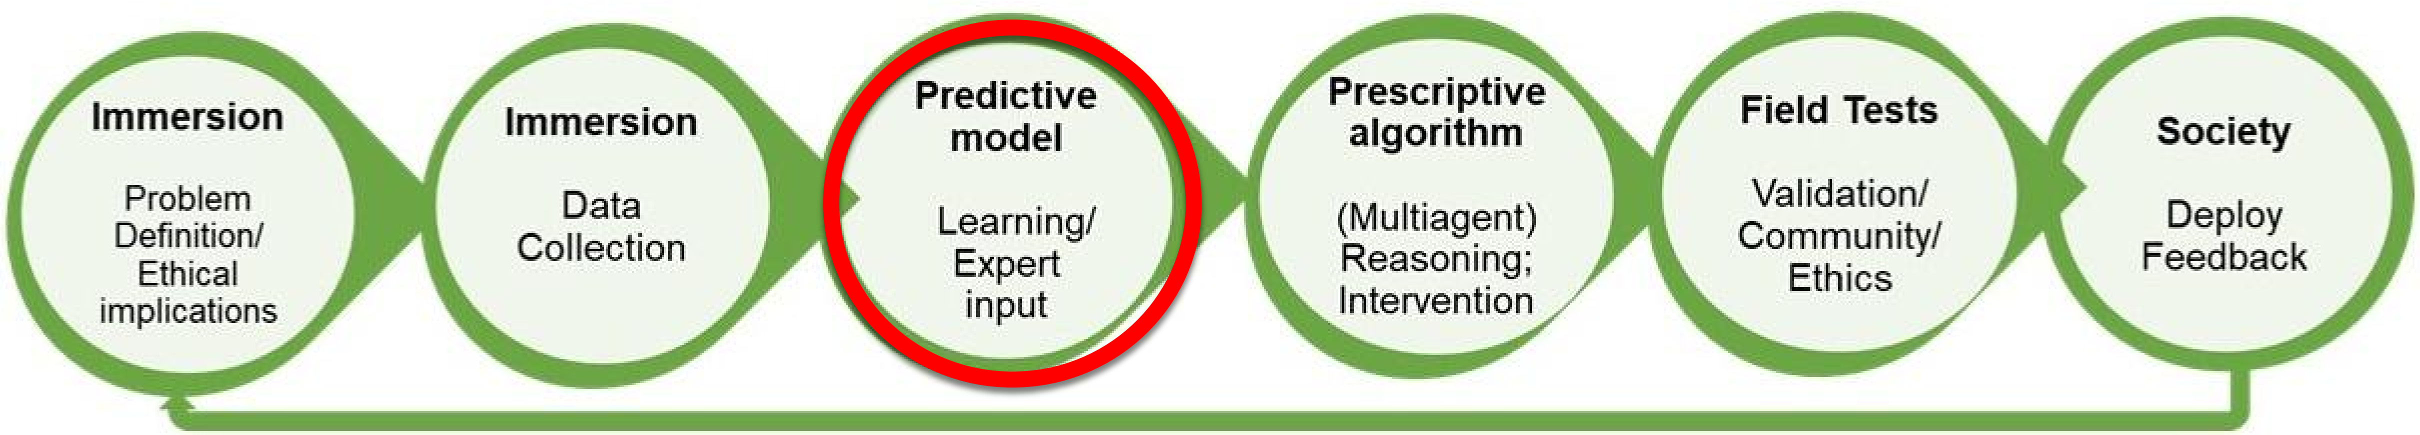

Key Lessons for Today

- Fairness challenges extend beyond the prediction phase

- Popular fairness definitions have shortcomings

- How papers evolve: NeurIPS’19→Workshop→AAAI’21

Fairness in Predictive vs Prescriptive Phase

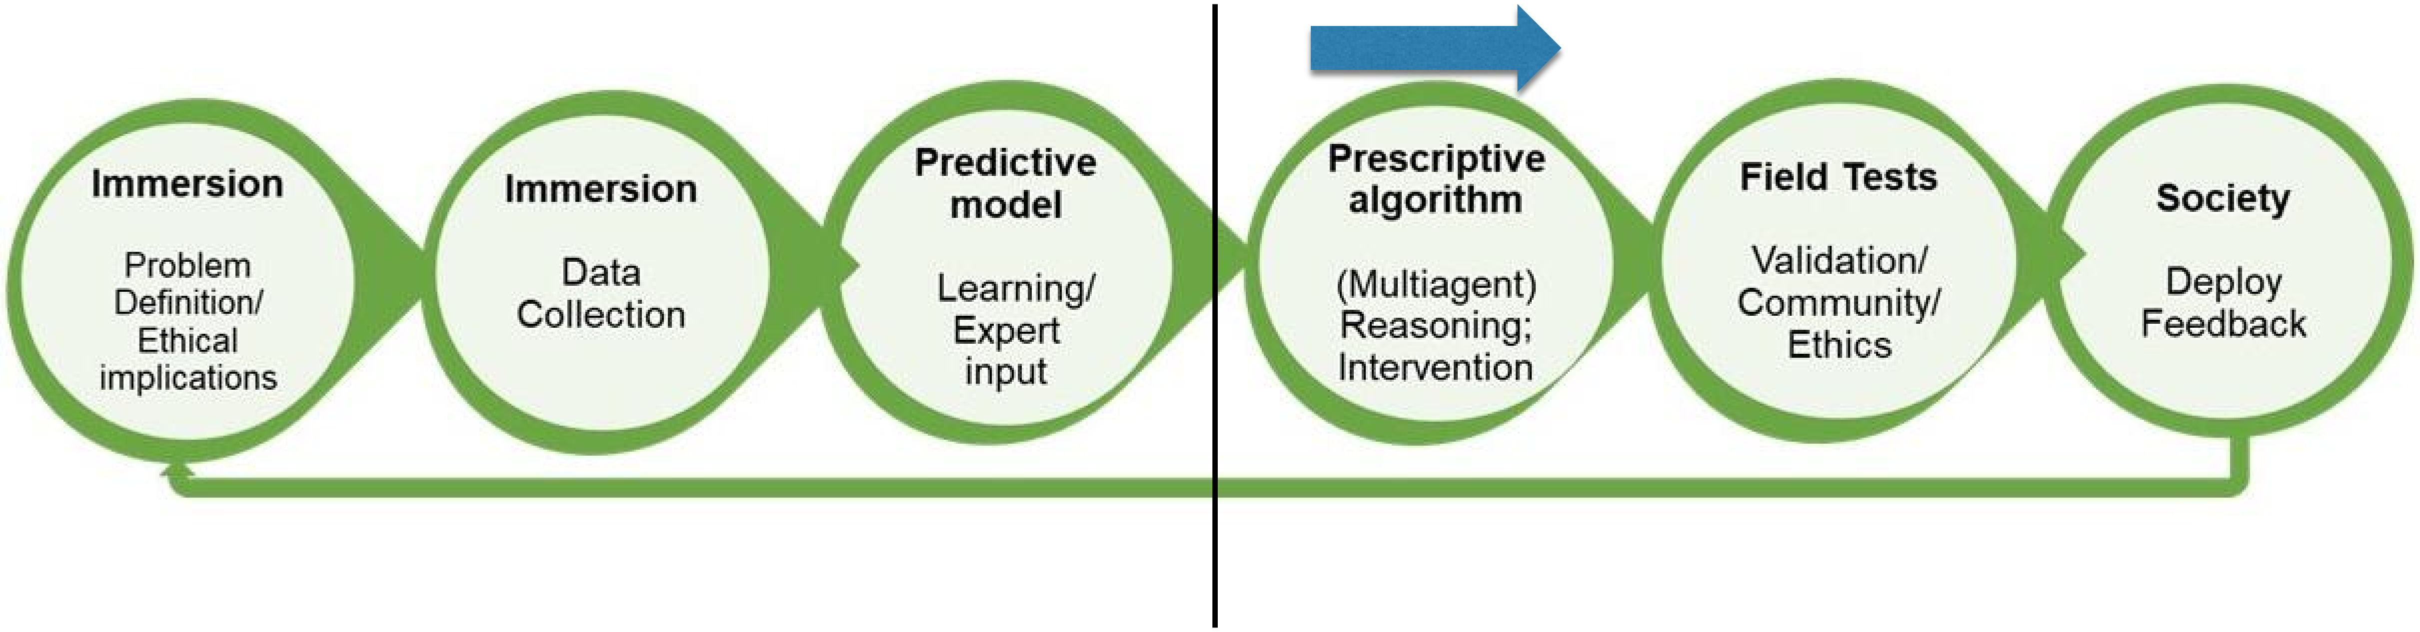

Algorithmic Fairness in Suicide Prevention Interventions for the Homeless Youth Population

Aida Rahmattalabi*, Phebe Vayanos, Anthony Fulginiti, Eric Rice, Bryan Wilder, Amulya Yadav, Milind Tambe

Alarming Rates of Suicide

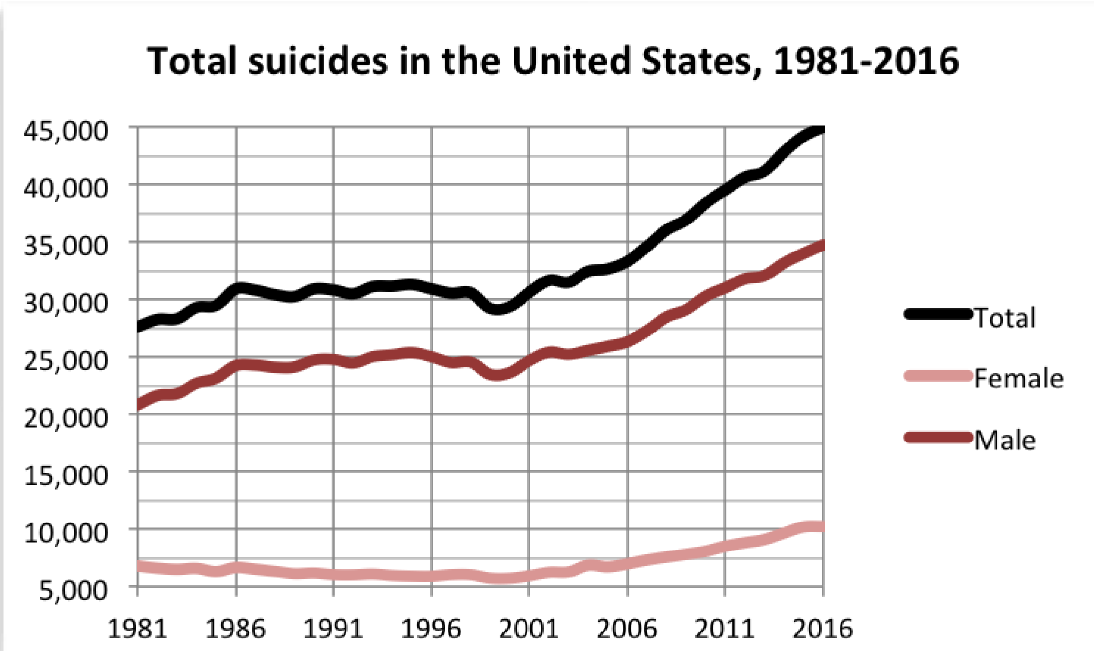

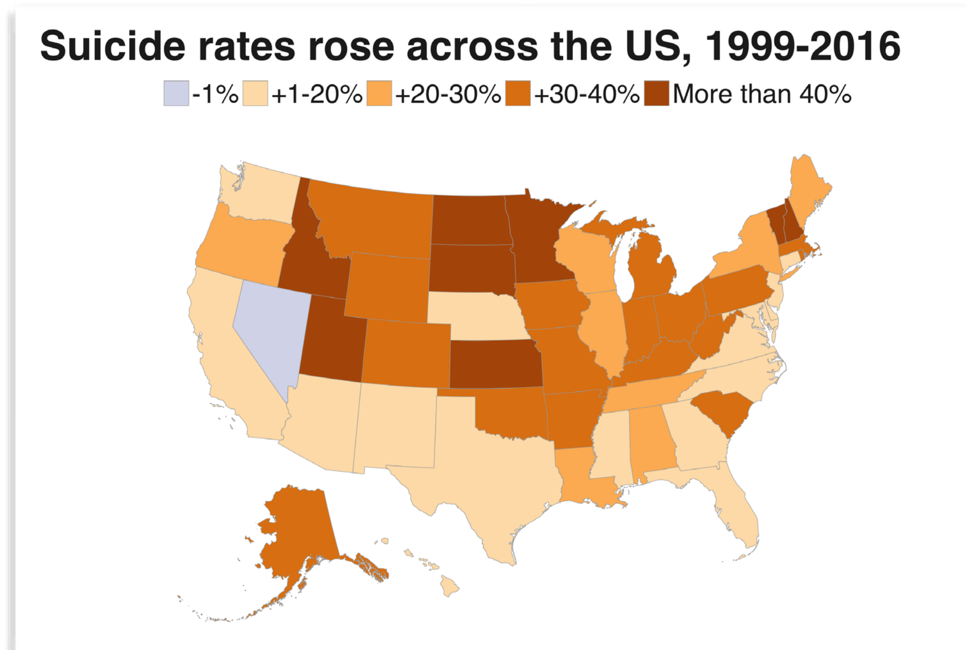

- 25% nationwide increase during 1999-2016
- Second leading cause of death among ages 10-34!

National Health Care for the Homeless Council, Fact Sheet, 2018

Suicide Rate and Homelessness

- Suicide rate is significantly high among youth experiencing homelessness (~7%-26%)

National Health Care for the Homeless Council, Fact Sheet, 2018

“Gatekeeper” Training

- One of the most popular suicide prevention program
- Conducted among college students, military personnel, etc.
- Trains “helpers” to identify warning signs of suicide and how to respond

Can we leverage
social network information?

Challenges

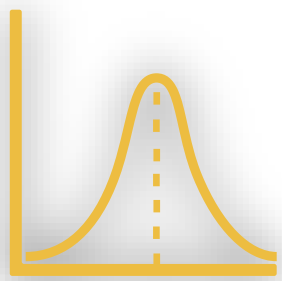

- Uncertainty in availability and performance

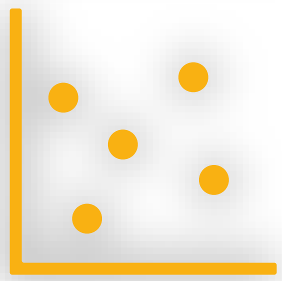

- Limited data to inform individuals' availability/performanc

- Combinatorial explosion in number of scenarios

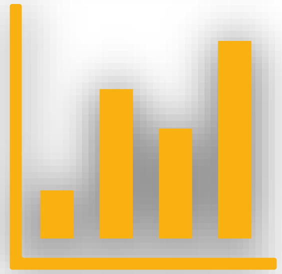

- Fairness in the delivery of the interventions

Intervention Model

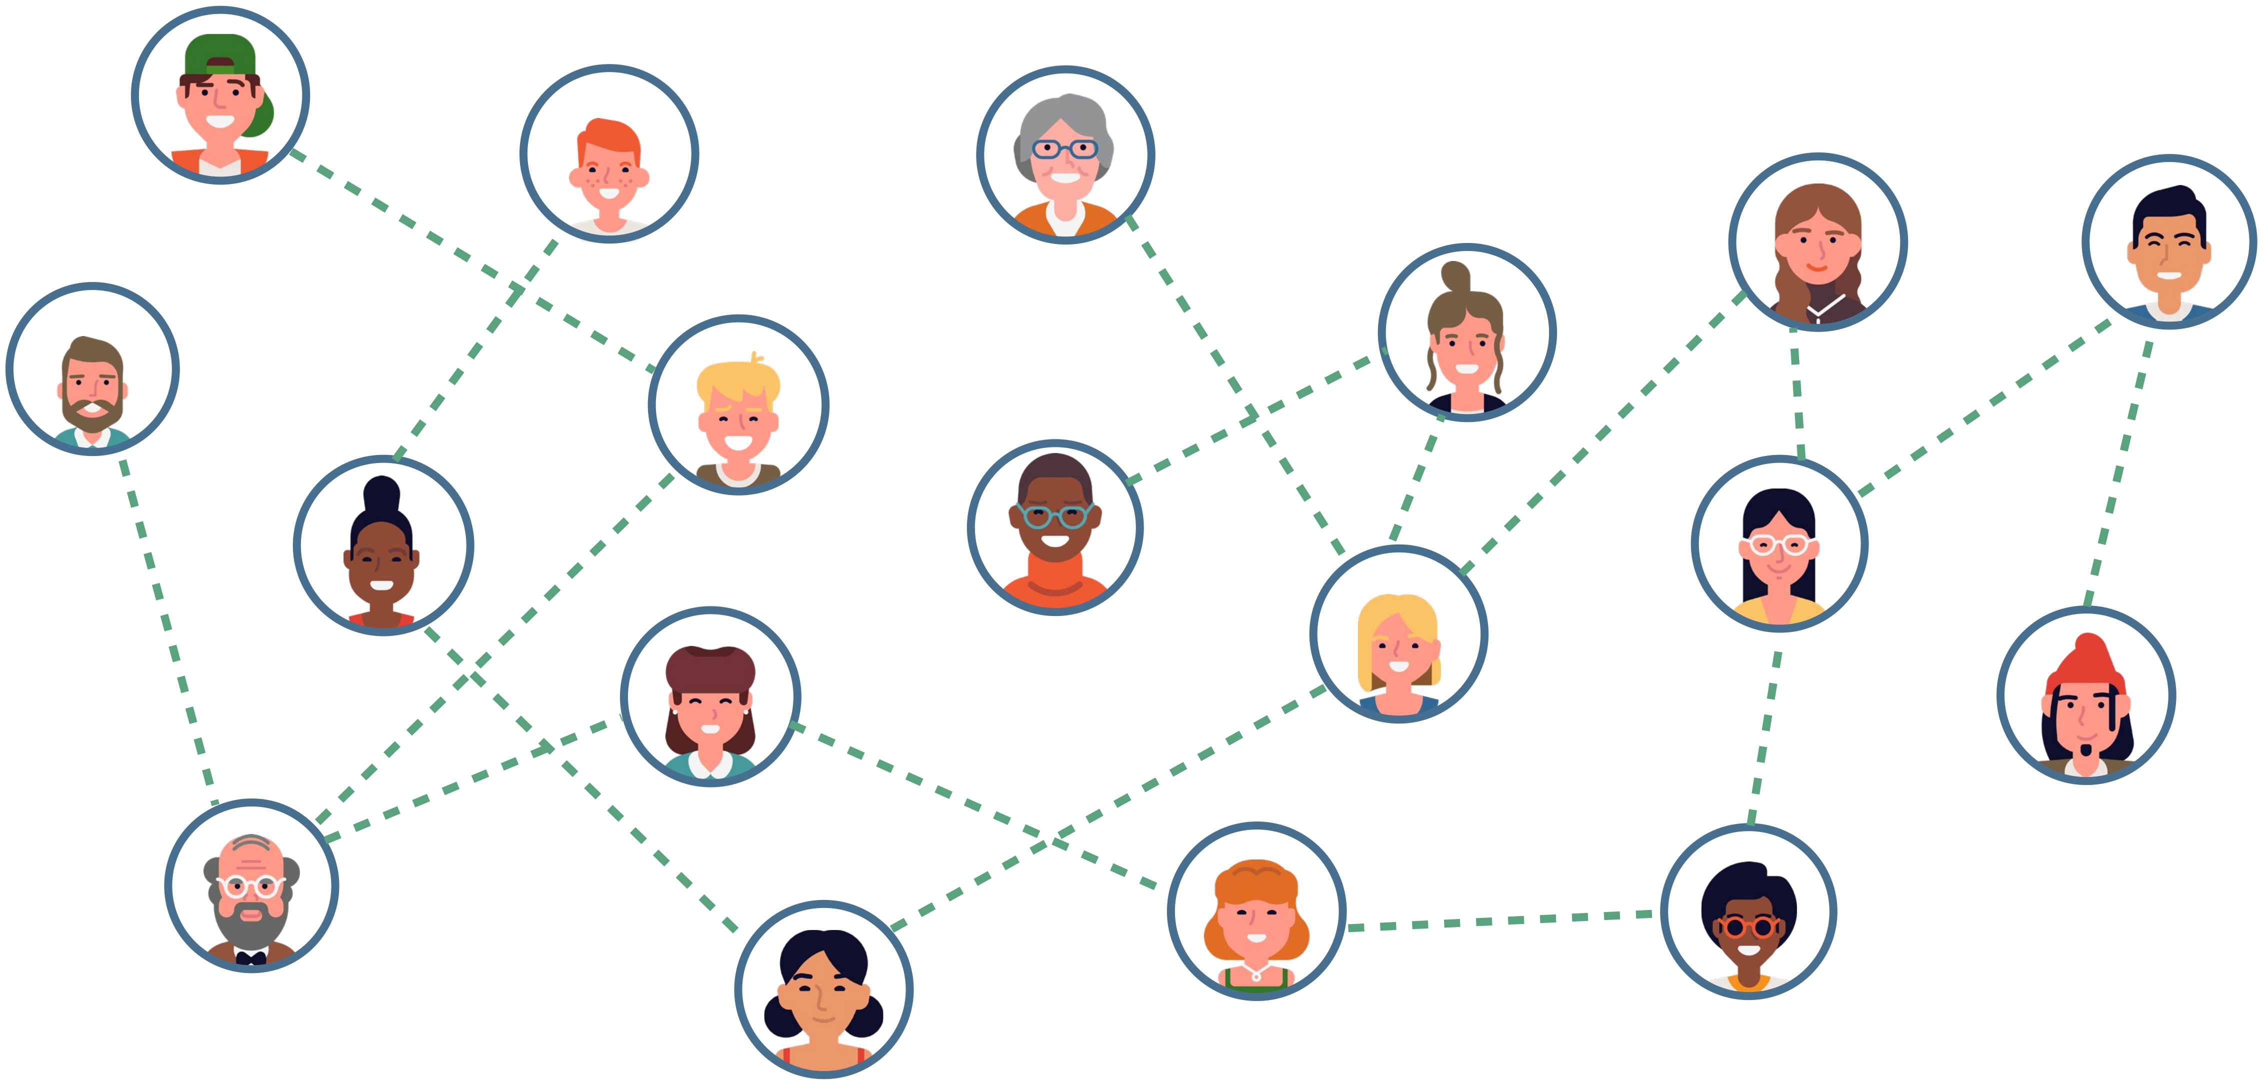

Known social network: $\mathcal{G} = (\mathcal{N}, \mathcal{E})$

Intervention Model

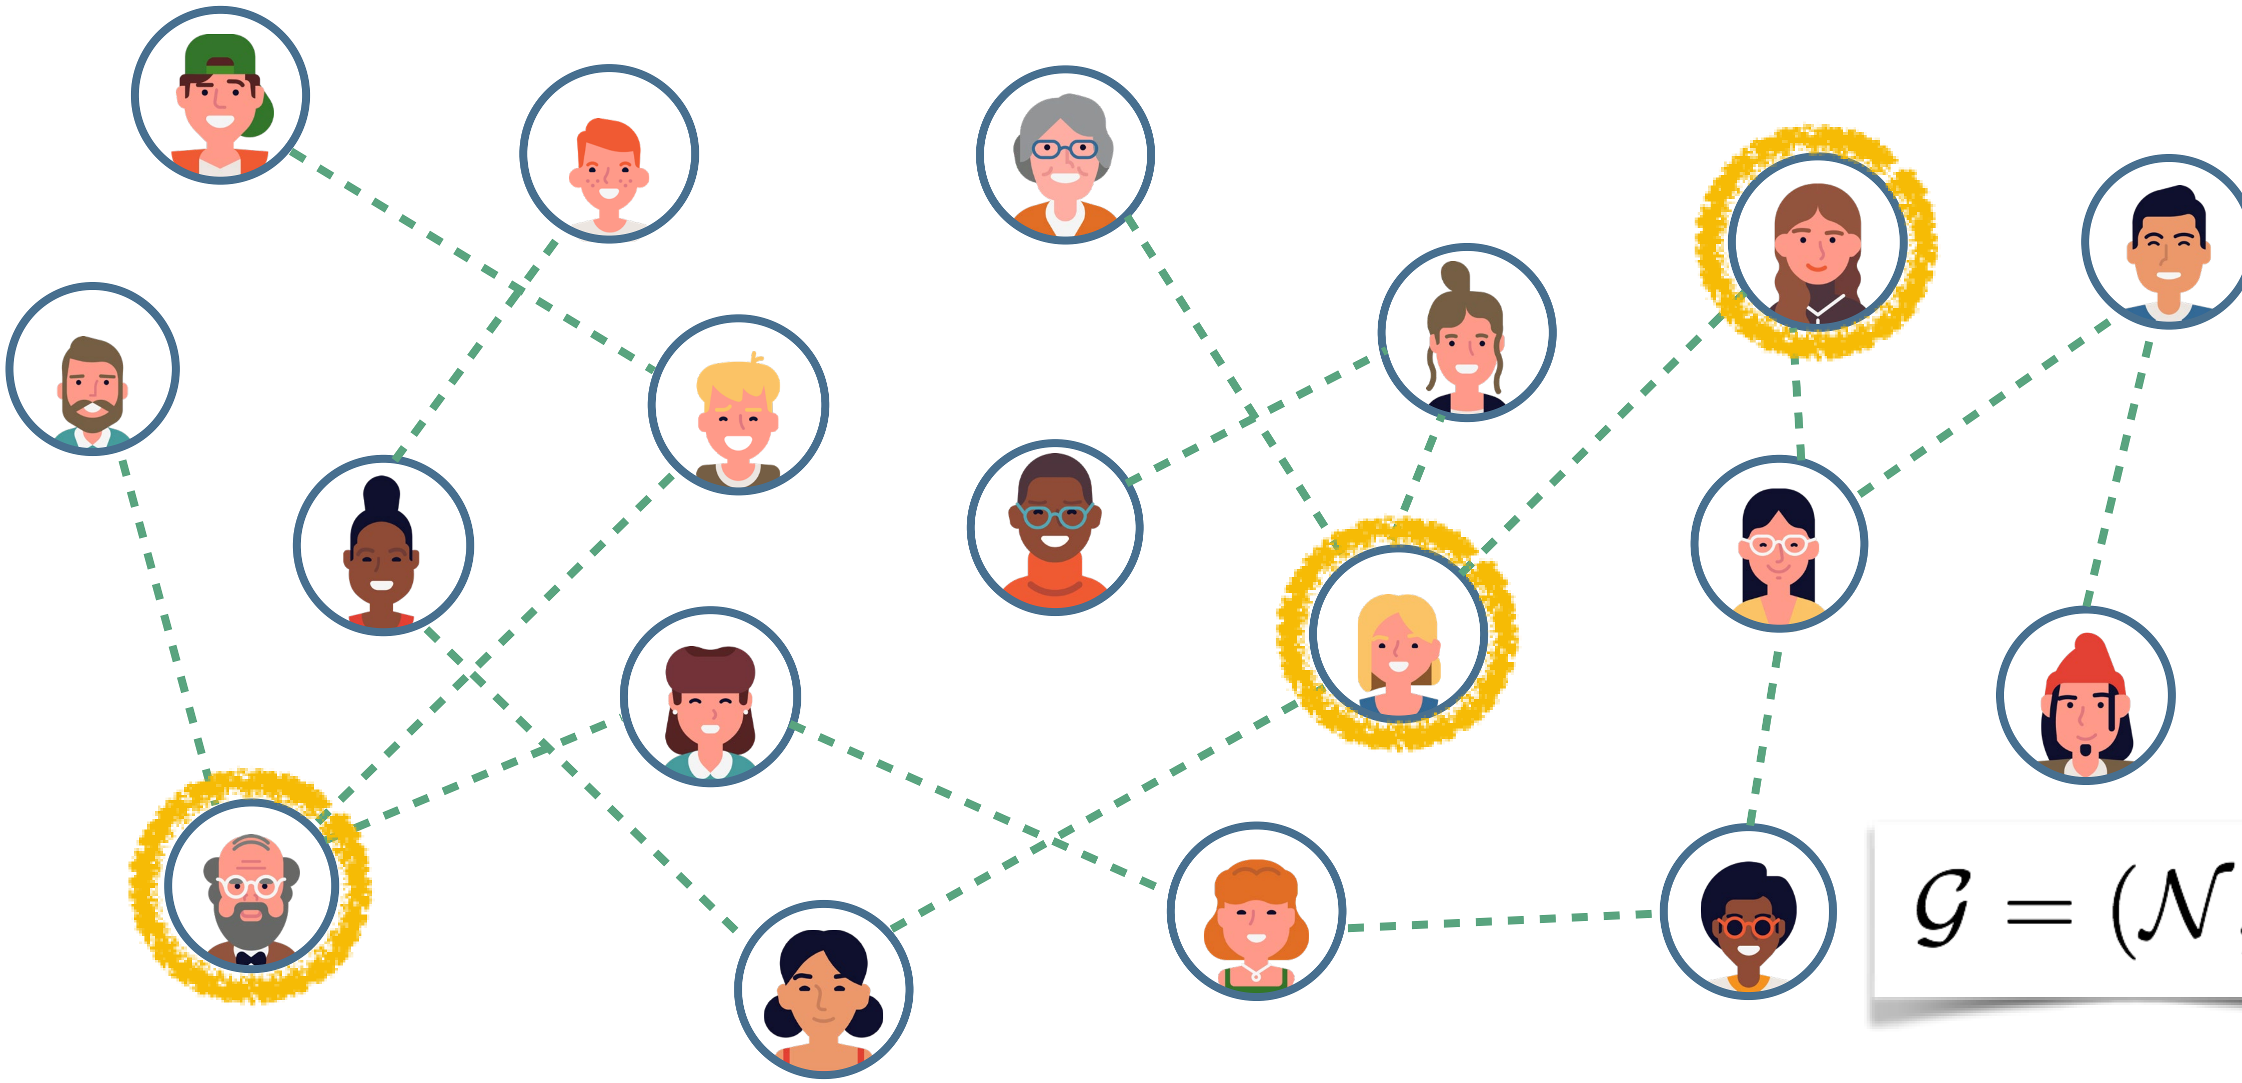

$$\mathcal{G} = (\mathcal{N}, \mathcal{E})$$

Train as a monitor:

$$x_n = 1$$

Intervention Model

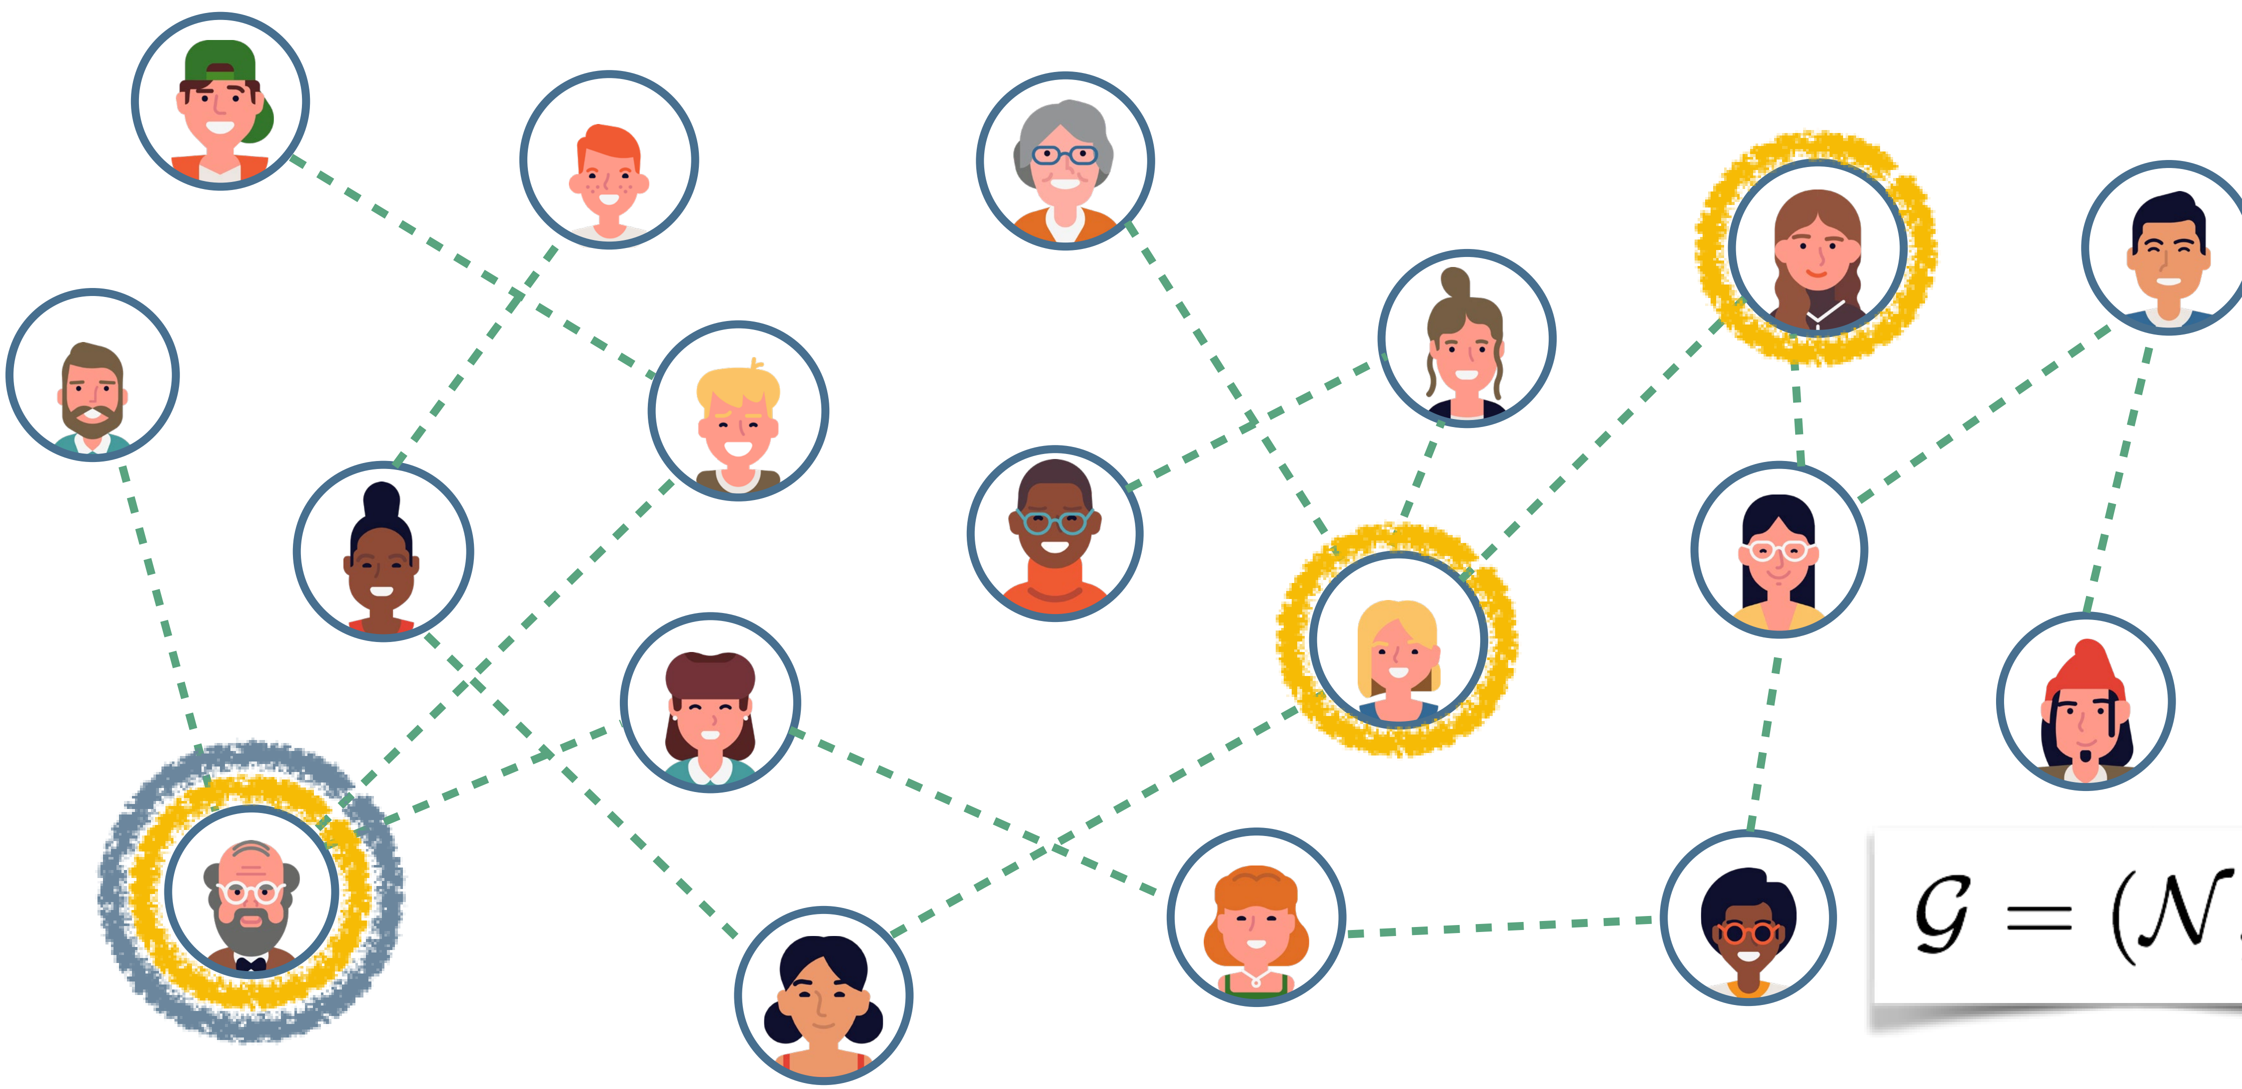

$\mathcal{G} = (\mathcal{N}, \mathcal{E})$

Train as a gatekeeper:

$x_n = 1$

Available:

$\xi_n = 1$

Intervention Model

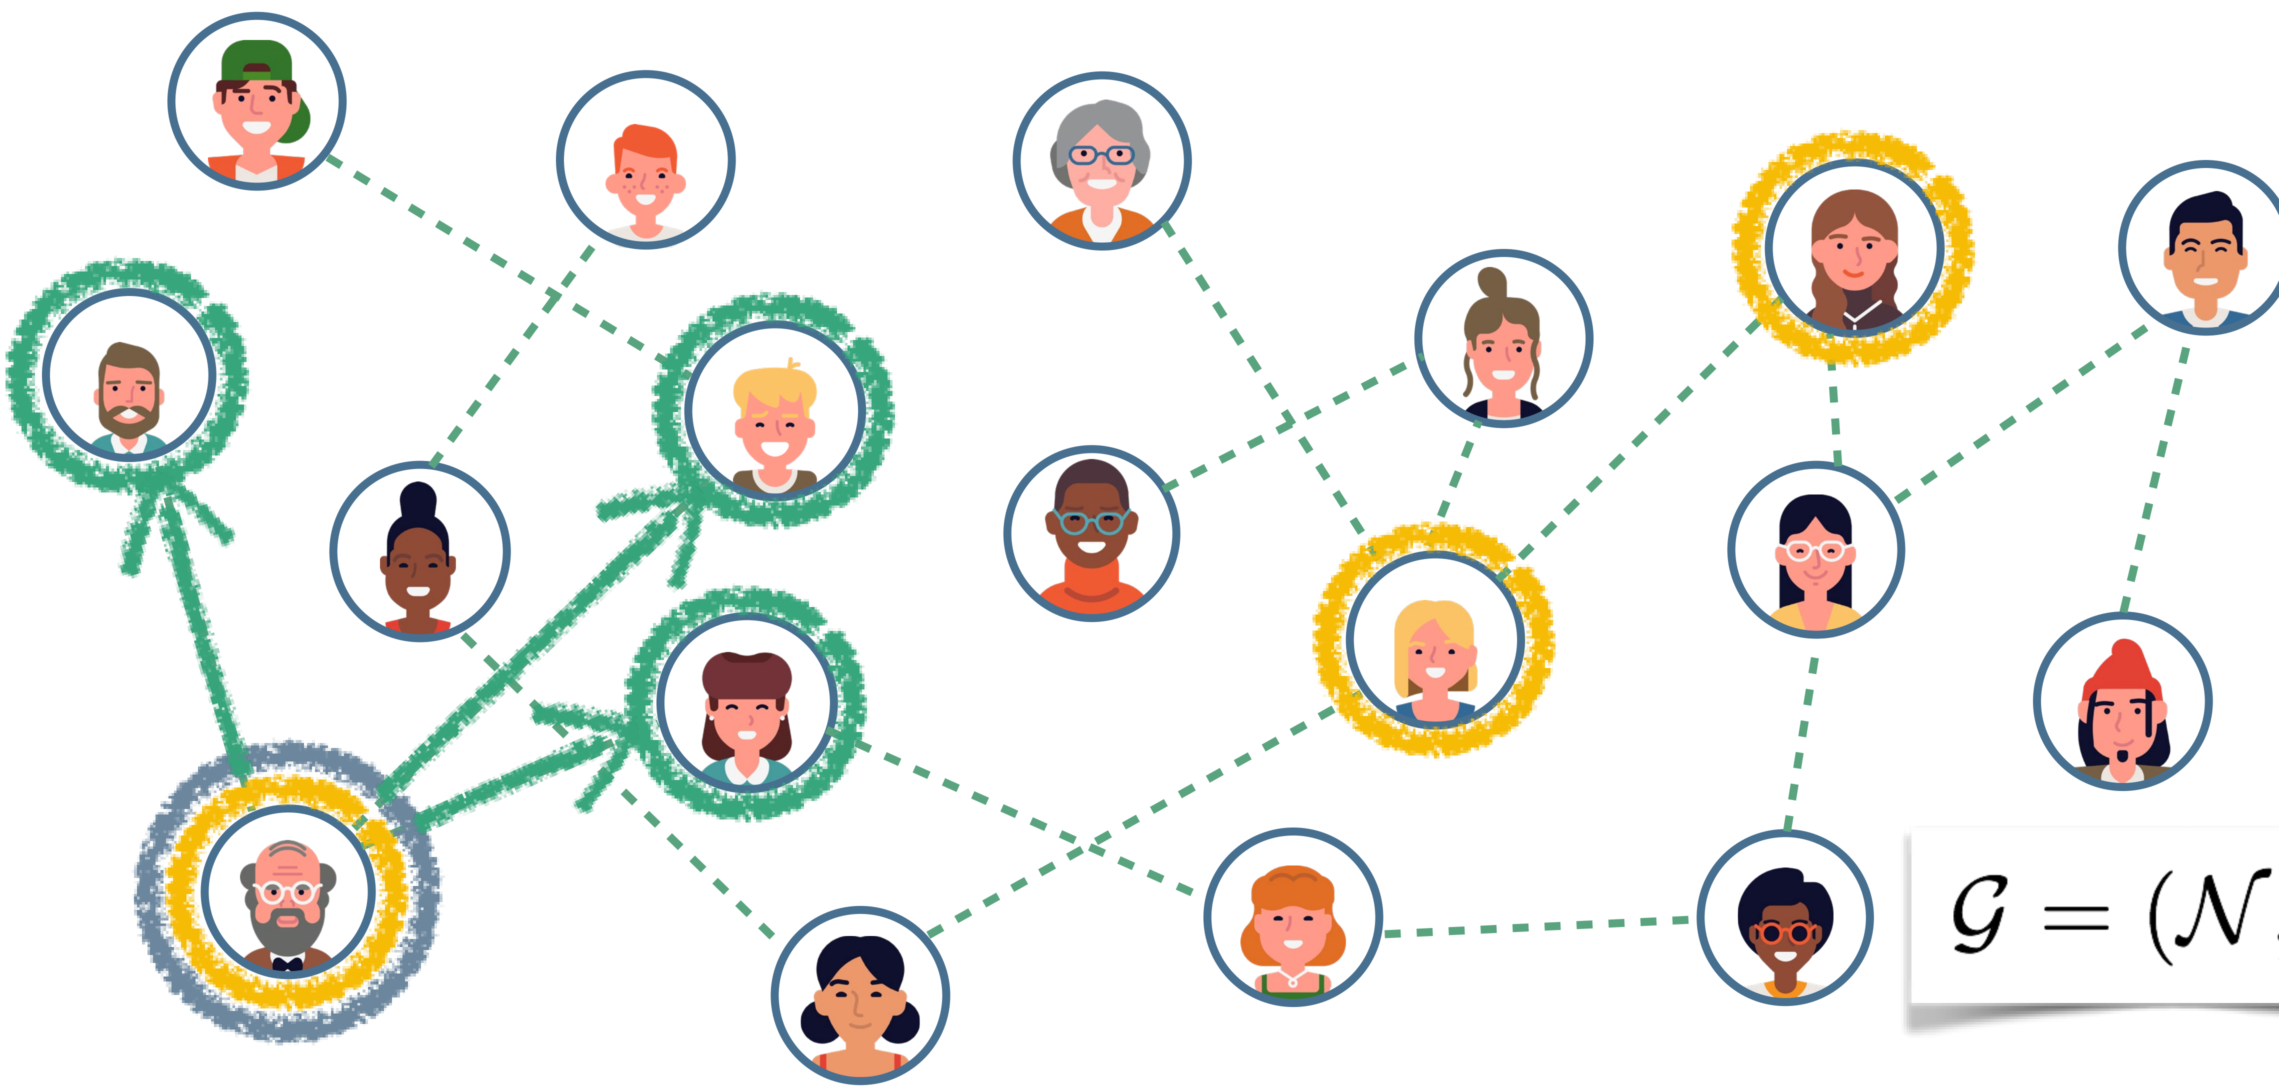

$$\mathcal{G} = (\mathcal{N}, \mathcal{E})$$

Train as a gatekeeper:

$x_n = 1$

Available:

$\xi_n = 1$

Covered:

$y_n(\boldsymbol{x}, \boldsymbol{\xi}) = 1$

Intervention Model
We maximize, nature minimizes

Robust Covering

$$\max_{\boldsymbol{x} \in \mathcal{X}} \min_{\boldsymbol{\xi} \in \Xi} F_{\mathcal{G}}(\boldsymbol{x}, \boldsymbol{\xi}) \text{ where } F_{\mathcal{G}}(\boldsymbol{x}, \boldsymbol{\xi}) := \sum_{n \in \mathcal{N}} y_n(\boldsymbol{x}, \boldsymbol{\xi})$$

Applying existing algorithm to Social Networks of Youth Experiencing Homelessness?

Suicide Prevention in Marginalized Populations: Choose Gatekeepers in social networks *(NeurIPS 2019)*

Rahmattalabi

Are we satisfied with these results?

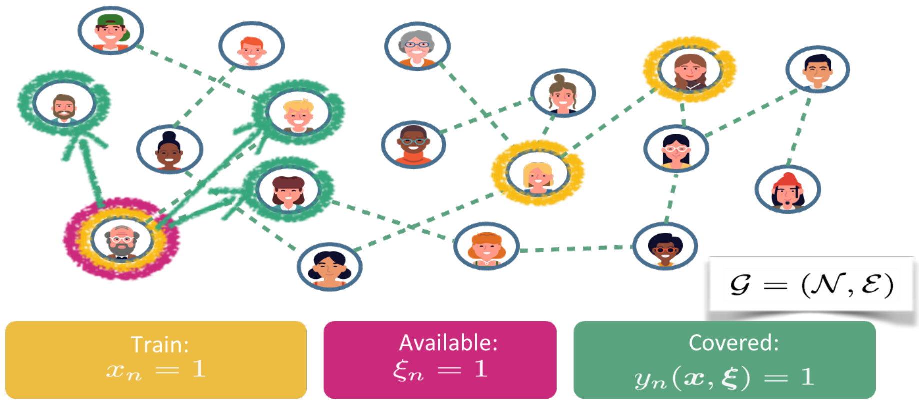

Network Name	Network Size	Percentage of covered individuals by racial group				
		White	Black	Hispanic	Mixed	Other
SPY1	95	70	36	–	86	94
SPY2	117	78	–	42	76	67
SPY3	118	88	–	33	95	69
MFP1	165	96	77	69	73	28
MFP2	182	44	85	70	77	72

Suicide Prevention in Marginalized Populations: Choose Gatekeepers in social networks *(NeurIPS 2019)*

Rahmattalabi

Is there any Inequality in coverage of marginalized populations?

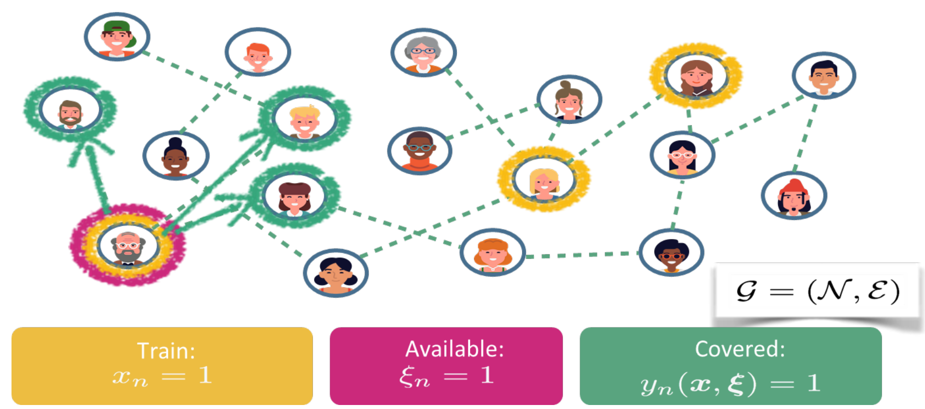

Network Name	Network Size	Percentage of covered individuals by racial group				
		White	Black	Hispanic	Mixed	Other
MFP1	165	96	77	69	73	28
MFP2	182	44	85	70	77	72

Suicide Prevention in Marginalized Populations: Choose Gatekeepers in social networks *(NeurIPS 2019)*

Rahmattalabi

- Inequality in coverage of marginalized population

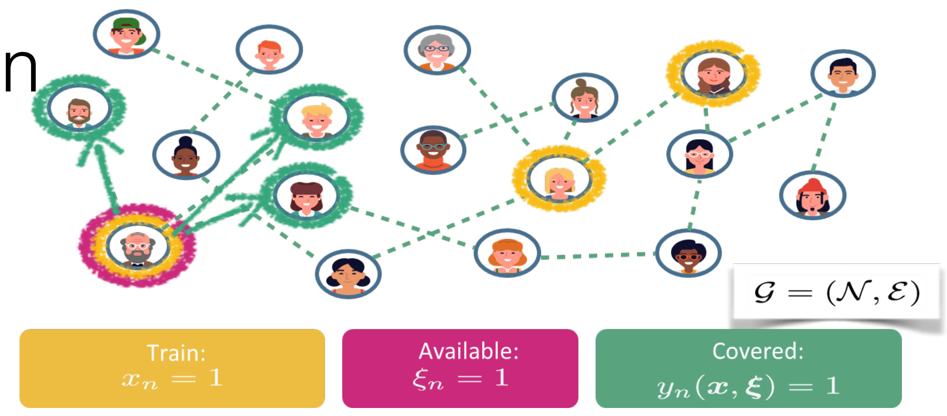

Network Name	Network Size	Percentage of covered individuals by racial group				
		White	Black	Hispanic	Mixed	Other
MFP1	165	96	77	69	73	28

Maximin Fairness:
Towards addressing Fairness Challenges

Rahmattalabi

- Maximin fairness constraints:
- Every group at least W (fraction) coverage

$$\max_{\boldsymbol{x}\in\mathcal{X}} \min_{\boldsymbol{\xi}\in\Xi} \max_{\boldsymbol{y}\in\mathcal{Y}} \left\{ \sum_{n\in\mathcal{N}_c} y_n : y_n \leq \sum_{\nu\in\delta(n)} \xi_\nu x_n, \; \forall n \in \mathcal{N}_c \right\}$$

$$\mathcal{Y} := \left\{ \boldsymbol{y} \in \{0,1\}^N : \sum_{n\in\mathcal{N}_c} y_n \geq W|\mathcal{N}_c| \quad \forall c \in \mathcal{C}, \; \forall \boldsymbol{\xi} \in \Xi \right\}$$

- Maximize W
- Can be done via binary search

Maximin Fairness: Towards addressing Fairness Challenges

Rahmattalabi

- Maximin fairness constraints: Every group at least W coverage

$$\max_{\boldsymbol{x}\in\mathcal{X}} \min_{\boldsymbol{\xi}\in\Xi} \max_{\boldsymbol{y}\in\mathcal{Y}} \left\{ \sum_{n\in\mathcal{N}_c} y_n : y_n \leq \sum_{\nu\in\delta(n)} \xi_\nu x_n, \ \forall n \in \mathcal{N}_c \right\}$$

$$\mathcal{Y} := \left\{ \boldsymbol{y} \in \{0,1\}^N : \sum_{n\in\mathcal{N}_c} y_n \geq W|\mathcal{N}_c| \quad \forall c \in \mathcal{C}, \ \forall \boldsymbol{\xi} \in \Xi \right\}$$

Network Name	Network Size	Percentage of covered individuals by racial group				
		White	Black	Hispanic	Mixed	Other
MFP1	165	85	69	42	73	64
MFP1	165	96	77	69	73	28

Are you happy with Maximin? Are there any issues you see in Maximin?

Maximin Fairness:
Towards addressing Fairness Challenges

Rahmattalabi

- Maximin fairness constraints: Every group at least W coverage

$$\max_{\boldsymbol{x}\in\mathcal{X}} \min_{\boldsymbol{\xi}\in\Xi} \max_{\boldsymbol{y}\in\mathcal{Y}} \left\{ \sum_{n\in\mathcal{N}_c} y_n : y_n \leq \sum_{\nu\in\delta(n)} \xi_\nu x_n, \ \forall n \in \mathcal{N}_c \right\}$$

$$\mathcal{Y} := \left\{ \boldsymbol{y} \in \{0,1\}^N : \sum_{n\in\mathcal{N}_c} y_n \geq W|\mathcal{N}_c| \quad \forall c \in \mathcal{C}, \ \forall \boldsymbol{\xi} \in \Xi \right\}$$

Network Name	Network Size	Percentage of covered individuals by racial group				
		White	Black	Hispanic	Mixed	Other
MFP1	165	85	69	42	73	64
MFP1	165	96	77	69	73	28

What might be some properties of fairness algorithms that we might be interested in the context of social network interventions?
Should no group be worse off?
Should rich not get richer?

Fairness in Public Health Preventative Interventions

A. Rahmattalabi[*], S. Jabbari[†], P. Vayanos[*], H. Lakkaraju[†], M. Tambe[†]

[*] University of Southern California

[†] Harvard University

Summary

- Implications of popular definitions of fairness in the context of public health interventions
- Proposed desiderata to ensure fairness:
  - Should not exhibit Population separation and leveling down
- Seemingly simple desiderata not satisfied
- Future Work: New fairness definitions that do not suffer from these pitfalls

Public Health Interventions

- Allocation of scarce resources

Public Health Interventions

- Allocation of scarce resources
- Examples: HIV intervention or suicide prevention
  - Train a limited number of peer-leaders

Resource Allocation Problem

- Total population $\mathcal{N}$

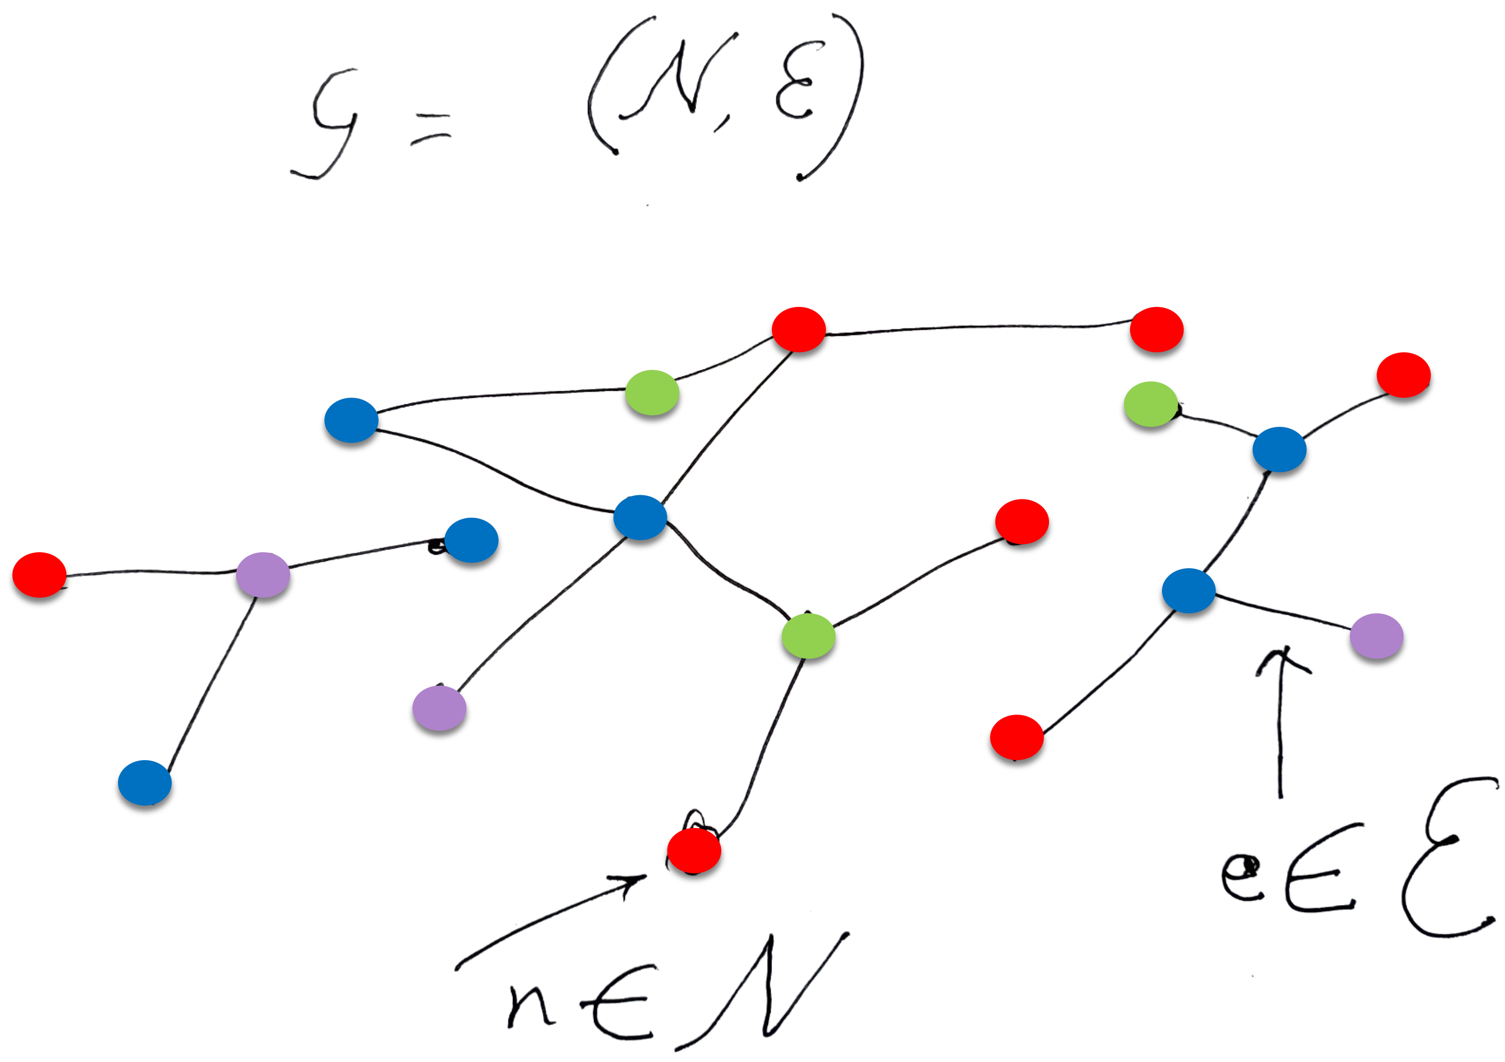

Resource Allocation Problem

- Total population $\mathcal{N}$
- Select peer leaders
- Maximize # of influenced nodes

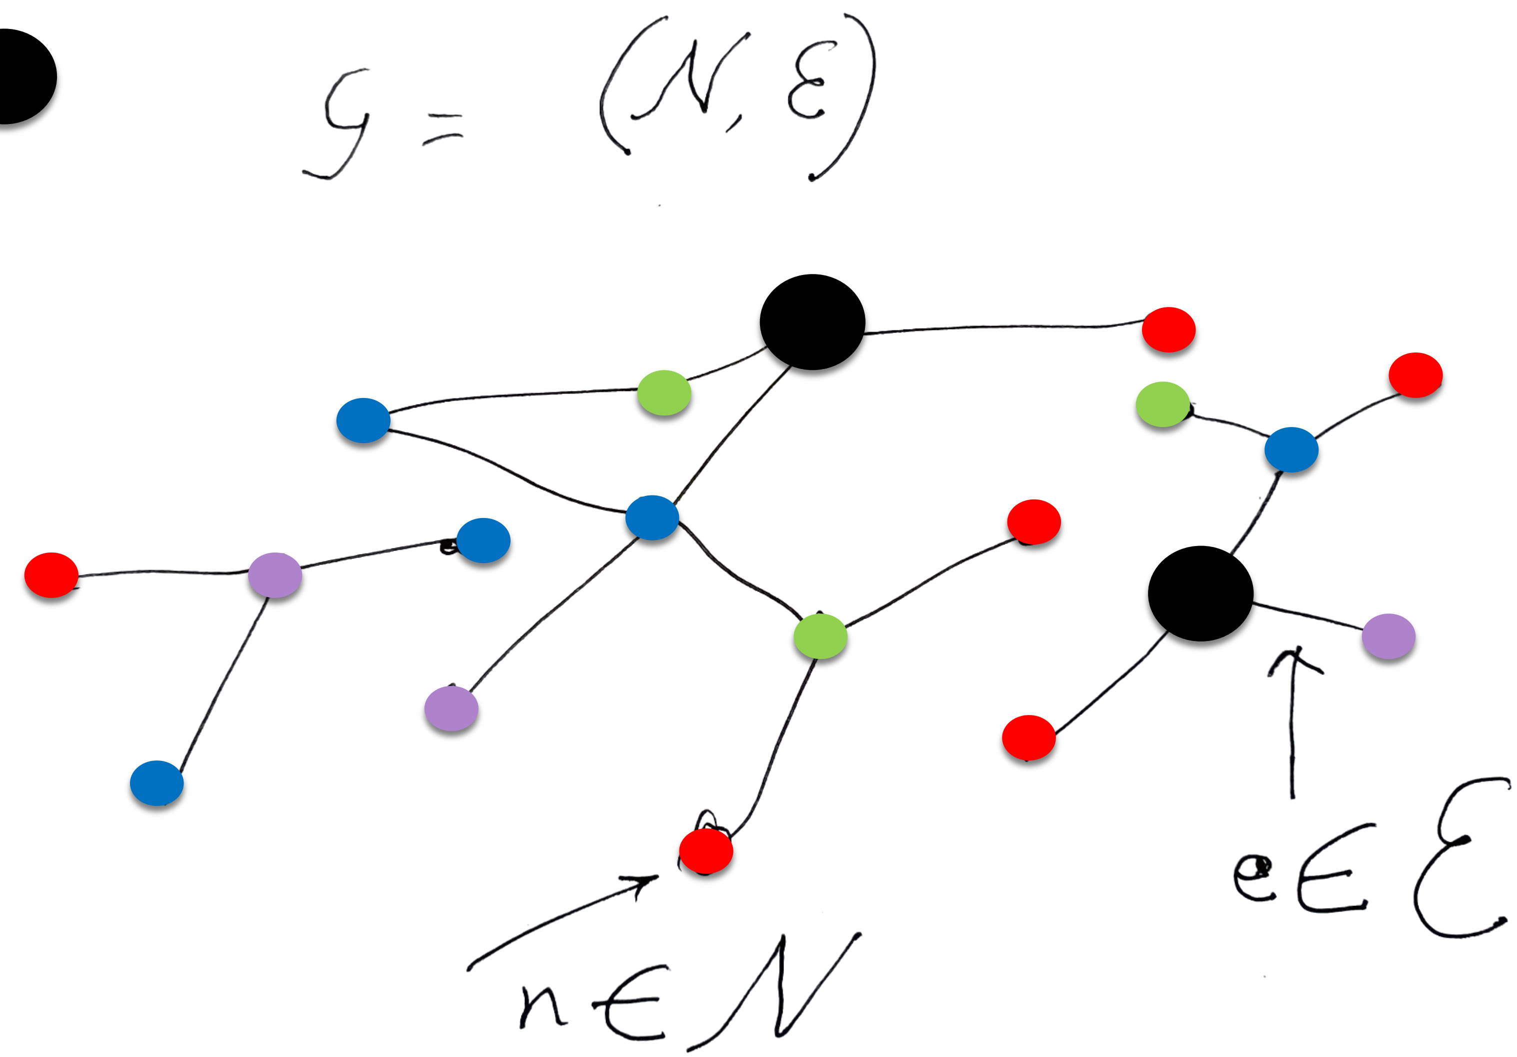

Resource Allocation Problem

- Total population $\mathcal{N}$
- Population consists of C disjoint groups $\mathcal{N}_1, \dots, \mathcal{N}_C$
- Utility of ● = (Number of influenced red nodes)/(Number of red nodes)

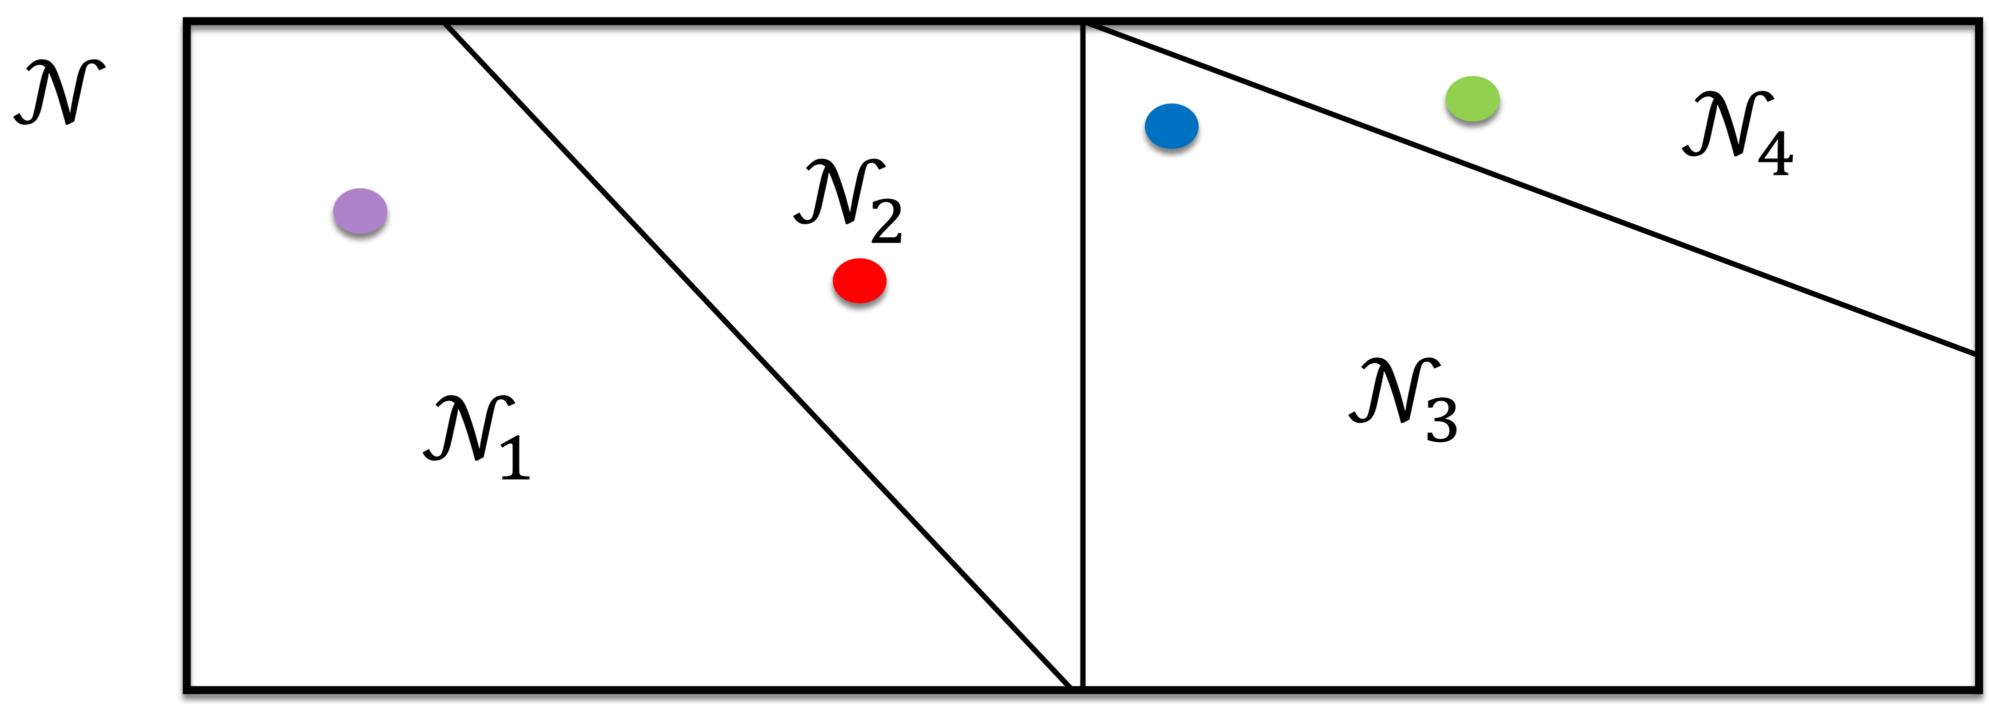

Resource Allocation Problem

- Total population $\mathcal{N}$
- Population consists of C disjoint groups $\mathcal{N}_1, \ldots, \mathcal{N}_C$
- Utility of 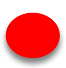 = (Number of influenced red nodes)/(Number of red nodes)
- utility of group "c", u_c = (Number of influenced nodes in "c")/$|\mathcal{N}_C|$ (e.g., 5/10 if 10 red nodes)

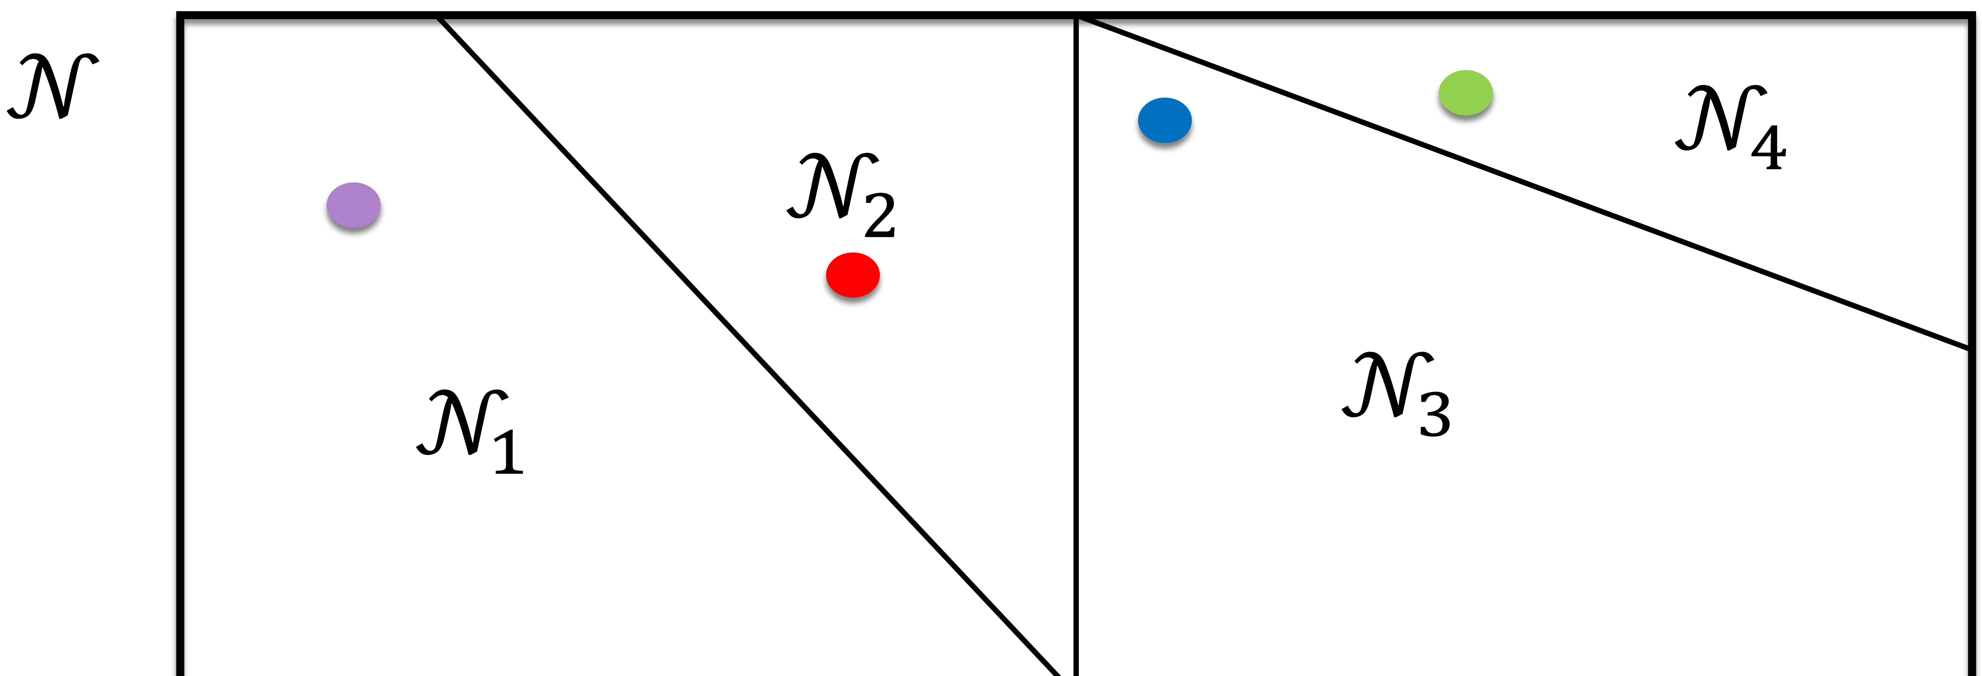

Objective of Resource Allocation Problem without Fairness

- C disjoint groups $\mathcal{N}_1, \ldots, \mathcal{N}_C$
- Without fairness considerations
- u_c = (Number of influenced nodes in "c")/| $\mathcal{N}_C$ |

$$\max_{\boldsymbol{u}} \sum_{c \in \mathcal{C}} |\mathcal{N}_c| \boldsymbol{u}_c : \boldsymbol{u} \in \mathcal{U}.$$

- Total population $\mathcal{N}$
- Select peer leaders ●
- Maximize # of influenced nodes

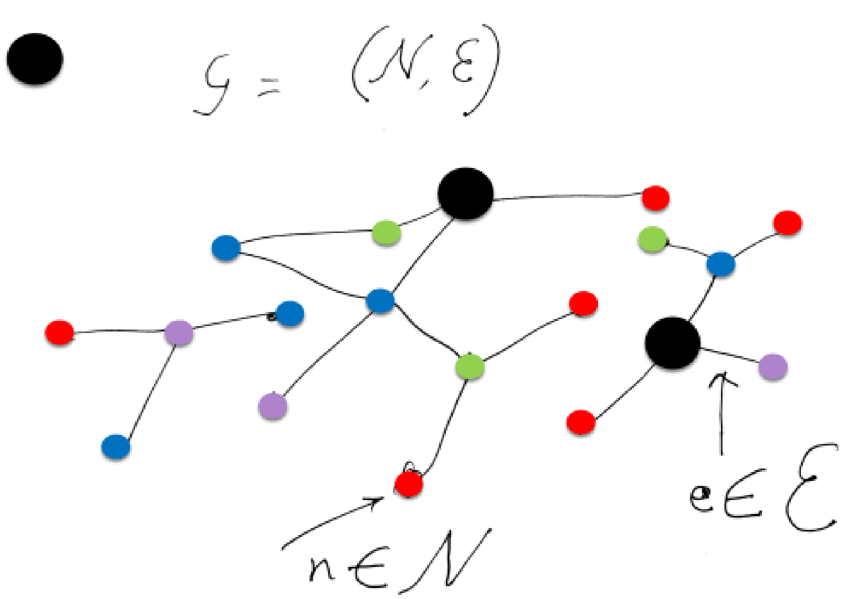

- U denotes possible utilities

Objective of Resource Allocation Problem with Fairness

- With fairness constraints on utilities

$$\max_{\boldsymbol{u}} \sum_{c \in \mathcal{C}} |\mathcal{N}_c| \boldsymbol{u}_c : \ \boldsymbol{u} \in \mathcal{U}, \boldsymbol{u} \in \mathcal{F}.$$

Fairness Definitions

- Maximin Fairness (MMF)

Maximize the utility of the worse-off group

Definition 1 (Maximin Fairness [15]). *The constraint set* $\mathcal{F}$ *for* maximin fairness (MMF) *is defined as*

$$\mathcal{F} = \left\{ \boldsymbol{u} \in [0,1]^C : \ \boldsymbol{u}_c \geq W \ \ \forall c \in \mathcal{C} \right\},$$

where $W := \sup\{v \in \mathbb{R}_+ \mid \exists \boldsymbol{u} \in \mathcal{U} : \boldsymbol{u}_c \geq v \ \ \forall c \in \mathcal{C}\}$ *is the maximum value for which Problem* (3) *remains feasible.*

Must be able to achieve at least W for each group

Fairness Definitions

- Maximin Fairness (MMF)

Maximize the utility of the worse-off group

- Demographic Parity (DP)

Every group must receive the same utility

Definition 2 (Demographic Parity). *Let* $\delta \in [0,1)$. *Then the constraint set* $\mathcal{F}$ *for* demographic parity (DP) *is defined as*

$$\mathcal{F} = \left\{ \boldsymbol{u} \in [0,1]^{C} : \boldsymbol{u}_{c'} - \boldsymbol{u}_{c} \leq \delta \;\; \forall c, c' \in \mathcal{C} \right\}.$$

Utilities of different groups must be within "delta" of one another

Fairness Definitions

- Maximin Fairness (MMF)

Maximize the utility of the worse-off group

- Demographic Parity (DP)

Every group must receive the same utility

- Diversity Constraints (DC)

No group should be better off by taking their proportional allocation of resources and allocating them internally

Desiderata
Population Separation

By imposing fairness, the utility gap should not increase

- Exhibit population Separation

By imposing fairness, the utility gap increases

Desiderata

By imposing fairness, the utility gap should not increase

- Should not exhibit Population Separation

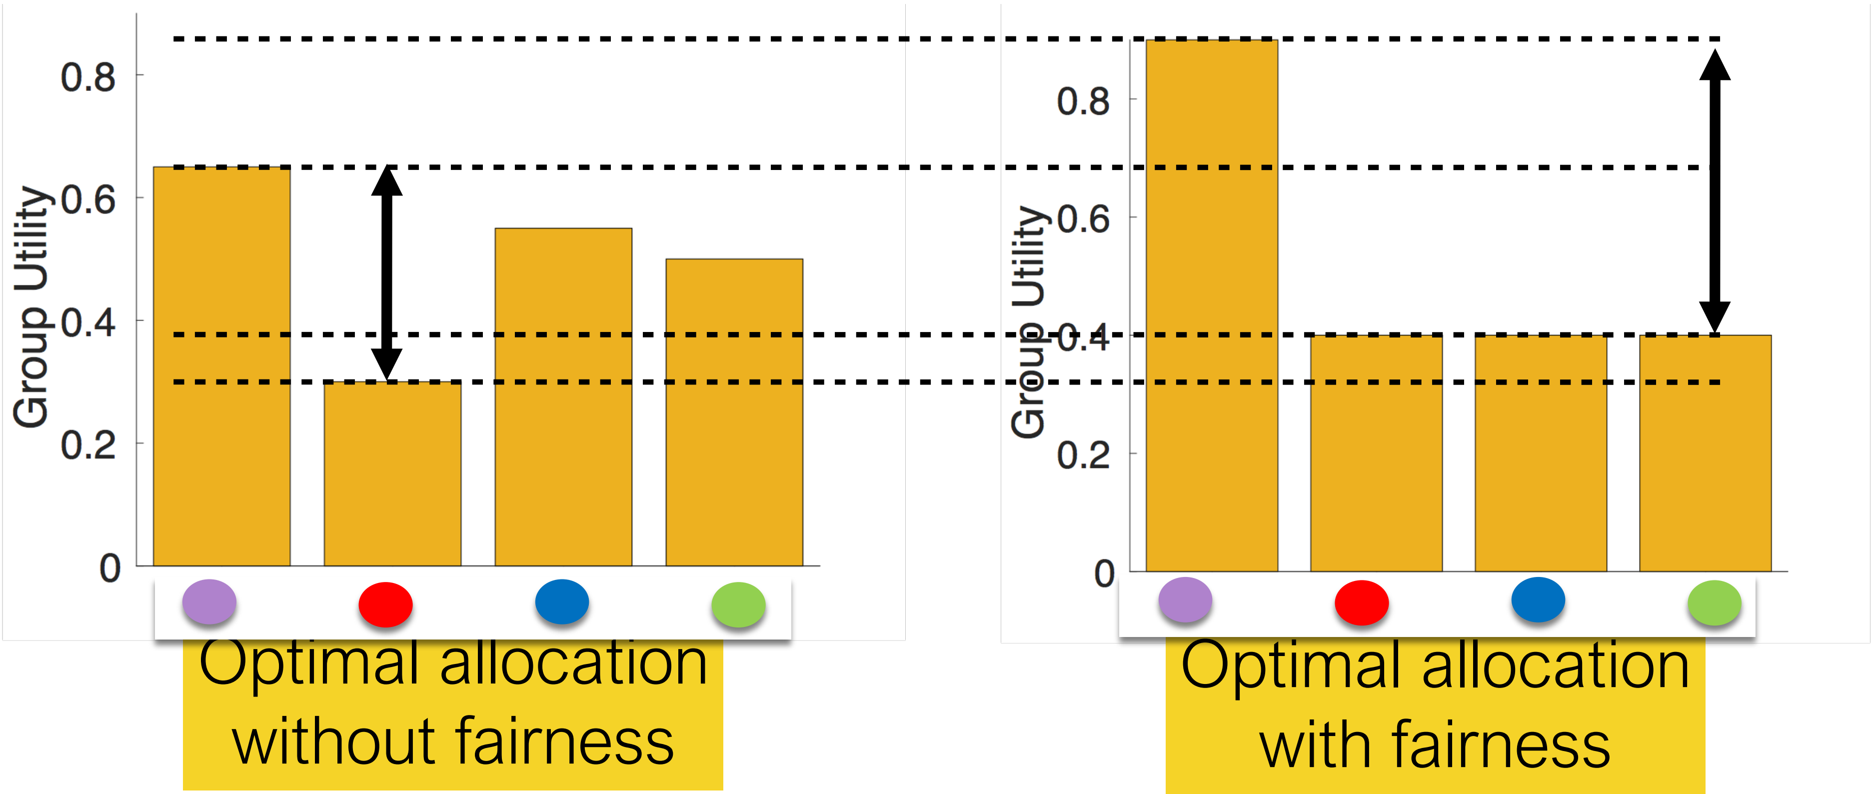

Desiderata
Leveling Down

By imposing fairness, the utility of all groups should not degrade

- Exhibit Leveling Down

By imposing fairness, the utility of all groups strictly degrades (strong leveling down)

By imposing fairness, the utility of no group increases and one group strictly degrades (weak leveling down)

Desiderata

By imposing fairness, the utility of all groups should not degrade

- Exhibits strong Leveling Down, but should not

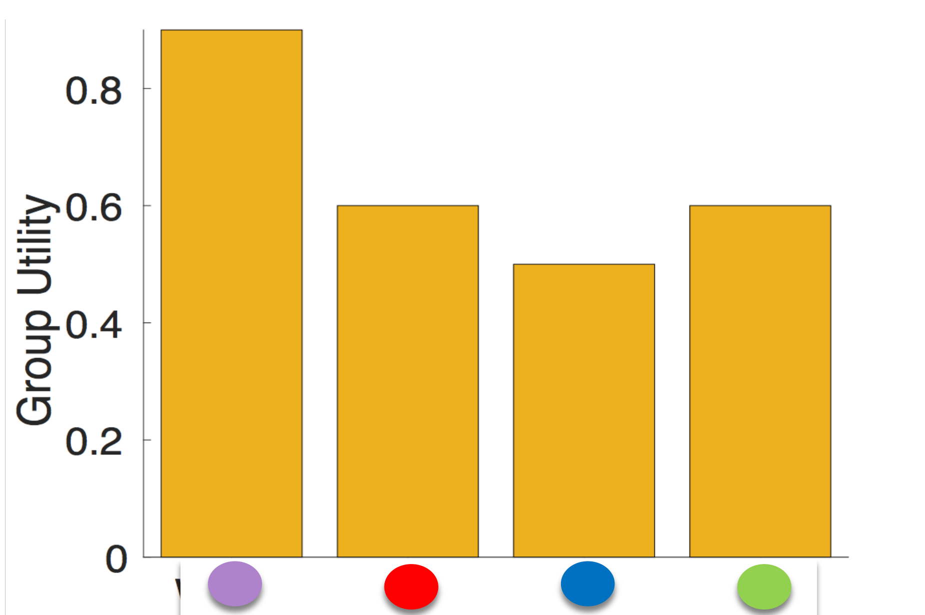

Optimal allocation without fairness

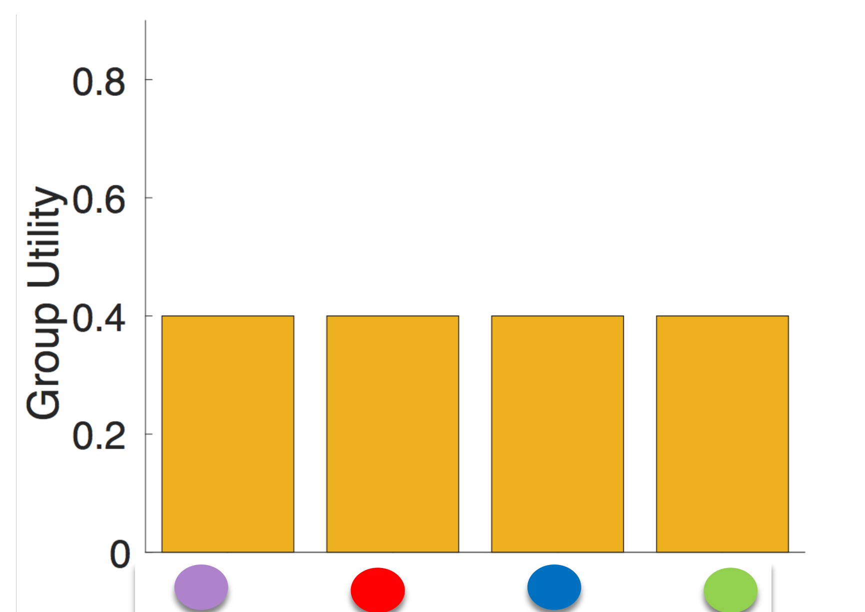

Optimal allocation with fairness

Do you agree? Do “your” proposed properties agree with the ones proposed here?

Theoretical Results

Maxmin Fairness & Population Separation property Group Exercise [PLEASE TAKE A SCREENSHOT]

One seed node

Optimal allocation without fairness:

Utilities of blue, red, green?

Optimal allocation with maxmin fairness?

Utilities of blue, red, green?

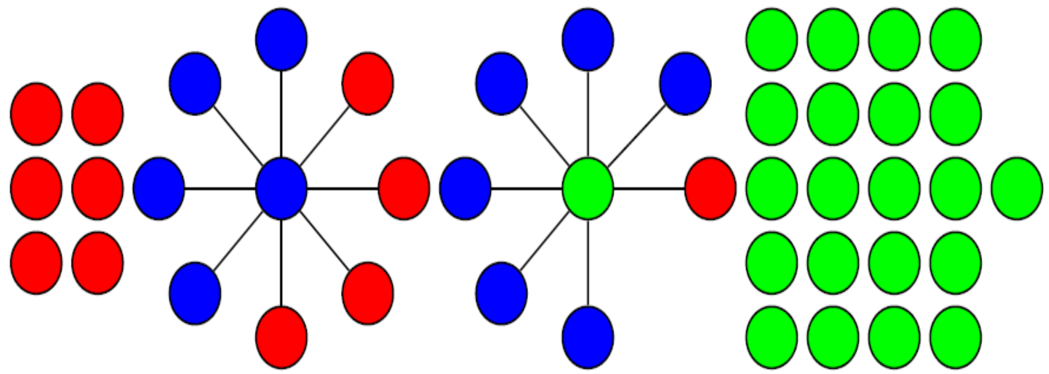

Figure 1: Companion figure to Proposition 1 for the case of $p = 1$. The network consists of three groups: red, blue and green. The edges are undirected so the influence can spread both ways. For arbitrary p, the number of isolated green vertices should scale to $\lceil 21/p \rceil$.

By imposing fairness, the utility gap should not increase!

But does it here?

Maxmin Fairness & Population Separation property

One seed node

Optimal allocation without fairness:

→ Center of bigger star

Utilities of blue, red, green?

→ 5/11, 4/11, 0

Optimal allocation with maxmin fairness?

→ Center of smaller star

Utilities of blue, red, green?

→ 6/11, 1/11, 1/22

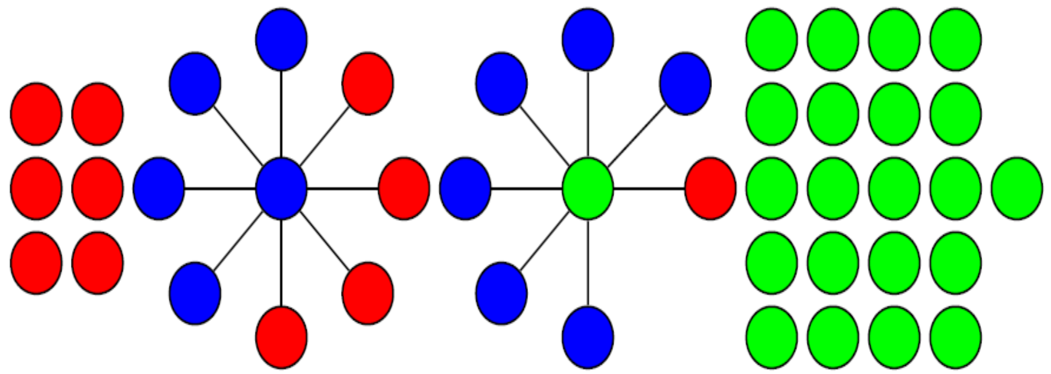

Figure 1: Companion figure to Proposition 1 for the case of $p = 1$. The network consists of three groups: red, blue and green. The edges are undirected so the influence can spread both ways. For arbitrary p, the number of isolated green vertices should scale to $\lceil 21/p \rceil$.

By imposing fairness, utility gap increases!

$$5/11 - 0 < 6/11 - 1/22$$

What if there were only two groups Red and Blue? Would Maxmin Fairness suffer from population separation?

Can Demographic Parity suffer from Population separation?

- Demographic Parity (DP)

Every group must receive the same utility

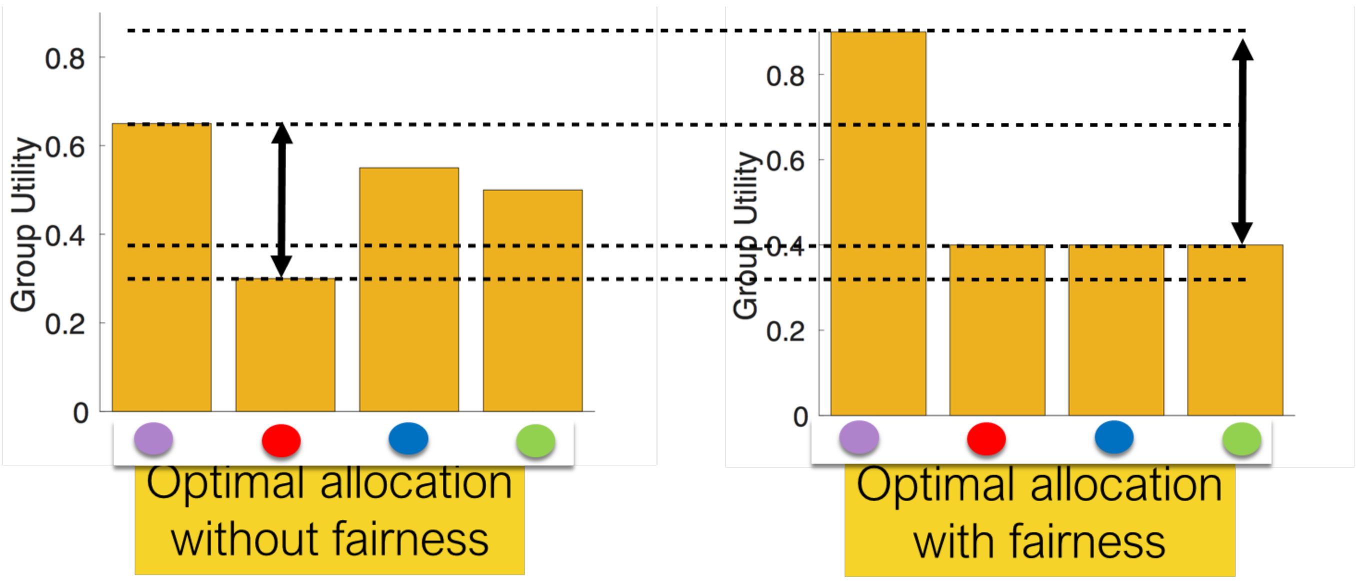

NO! Every group must receive the same utility in fair allocation (or within "delta")

Can Demographic Parity suffer from Leveling Down?

- Demographic Parity (DP)

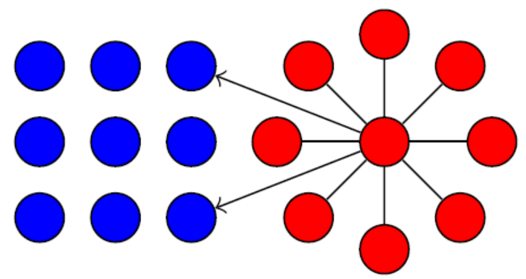

Figure 3: Companion figure to Proposition 5 where we choose $\max(3p/(p-\delta), (p-2p^2)/(\delta-p^2)) < N$, $\delta < p < \min(\sqrt{\delta}, 0.5)$, $k = 2$ and $T > 1$. The network consists of two groups red and blue, each of size N. All edges except the two shown by arrows are undirected meaning that influence can spread both ways.

- Delta = 0.25, p = 0.5
- Two seeds

- Without fairness: optimal allocation?

- With fairness optimal allocation?

Can Demographic Parity suffer from Leveling Down?

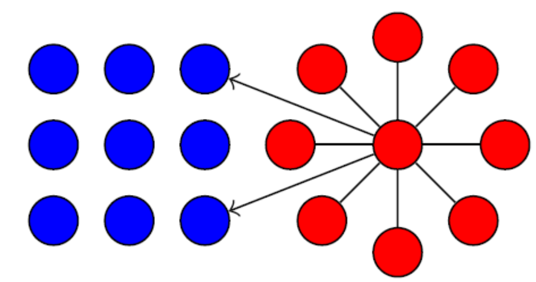

Figure 3: Companion figure to Proposition 5 where we choose $\max(3p/(p-\delta), (p-2p^2)/(\delta-p^2)) < N$, $\delta < p < \min(\sqrt{\delta}, 0.5)$, $k = 2$ and $T > 1$. The network consists of two groups red and blue, each of size N. All edges except the two shown by arrows are undirected meaning that influence can spread both ways.

- Demographic Parity (DP)

- Delta = 0.25, p = 0.5
- Two seeds

- Without fairness: optimal allocation? (Red: 5/9, Blue=2/9)
- Difference in utility is greater than 0.25
- With fairness optimal allocation? (Red: 3.25/9, Blue=1.5/9)
- Difference in utility is less than 0.25

IS THERE LEVELING DOWN IN THIS EXAMPLE?

- Demographic Parity (DP)

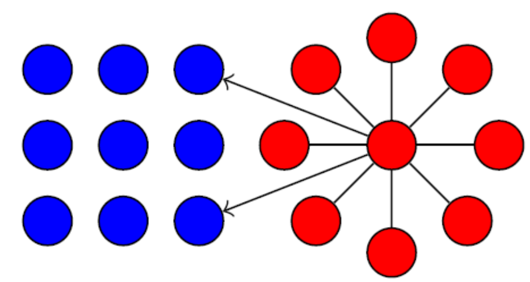

Figure 3: Companion figure to Proposition 5 where we choose $\max(3p/(p-\delta), (p-2p^2)/(\delta-p^2)) < N$, $\delta < p < \min(\sqrt{\delta}, 0.5)$, $k = 2$ and $T > 1$. The network consists of two groups red and blue, each of size N. All edges except the two shown by arrows are undirected meaning that influence can spread both ways.

- Delta = 0.25, p = 0.5
- Two seeds

- Without fairness: optimal allocation? (Red: 5/9, Blue=2/9)
- Difference in utility is greater than 0.25
- With fairness optimal allocation? (Red: 3.25/9, Blue=1.5/9)
- Difference in utility is less than 0.25

Theoretical Results

- MMF and DC both suffer from population separation
- DP suffers from leveling down

Fairness Criteria	Fairness Notion		
	MMF	DC	DP
Population Separation	✓	✓	✗
Leveling Down	✗	✗	✓

Theoretical Results

- Combination of fairness constraints suffers from the leveling down

Fairness Criteria	Fairness Notion		
	MMF+DP	DC+DP	DC+MMF+DP
Population Separation	✗	✗	✗
Leveling Down	✓	✓	✓

Conclusion and Future Work

- Implications of popular fairness definitions

- Proposed desiderata to ensure fairness:
    - Population separation and leveling down

- Future Work: New fairness definitions that do not suffer from these pitfalls

Fairness in Reasoning with Social Networks: Suicide Prevention via Gatekeeper Selection

(AAAI 2021)

Rahmattalabi

Choose K gatekeepers: robust & fair graph covering

Maximize worst case coverage with node failures & fairness constraints

Fairness in communities $c \in C$, allocations A

Maxmin fairness: $\min_{c \in C} u_c(A) \geq \gamma$ — γ: Max of minimum utility for any community

Diversity constraints: $u_c(A) \geq U_c$ — U_c: *Utility if # gatekeepers allocated proportional to size of community*

***Lesson**: Not commit to single fairness measure; each suffers from shortcomings*

but welfare maximization: $W_\alpha\left(u(A)\right)$

Allows adjusting tradeoffs by tuning "inequity aversion" parameter α

THANK YOU

A. Rahmattalabi[*], S. Jabbari[†], P. Vayanos[*], H. Lakkaraju[†], M. Tambe[†]

[*] University of Southern California, Center for Artificial Intelligence in Society

[†] Harvard University

Group-Fairness in Influence Maximization

Bryan Wilder[1]

Joint work with Alan Tsang[2], Yair Zick[2], Eric Rice[3], Milind Tambe[1]

[1]Harvard University

[2]National University of Singapore

[3]University of Southern California

Network interventions

- Disseminate information or change behavior via **social influence**
  - “Influence maximization”
- Often used in health and development
  - HIV prevention for homeless youth

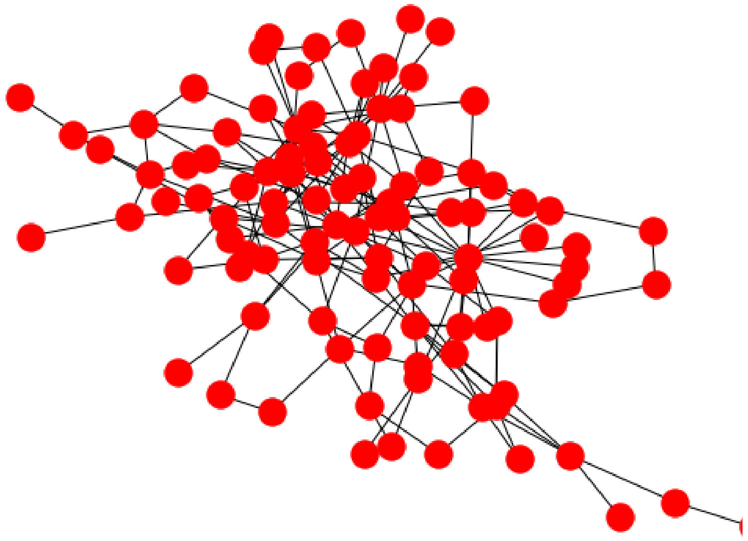

Fairness

- Typically, algorithms optimize **total welfare**
  - Expected number reached
- But how is influence distributed?
- Neglect small/less connected groups?

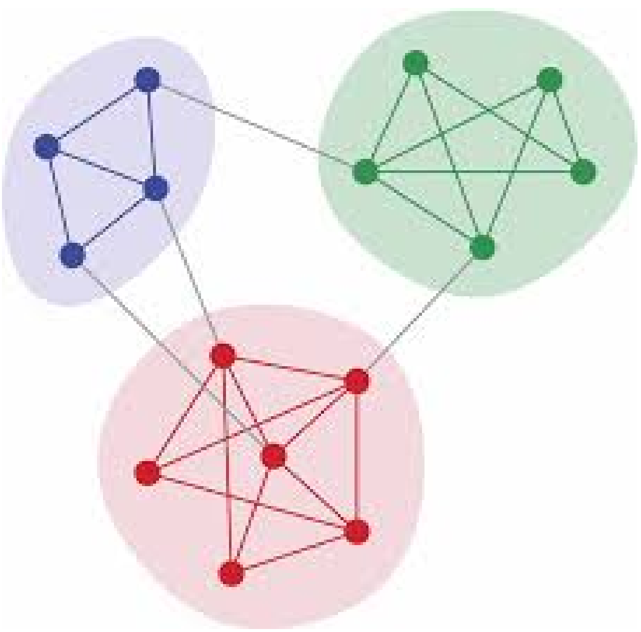

Do standard influence maximization algorithms neglect minority groups?

How can disparities be avoided?

Contributions

- Formalize measures of fairness in influence maximization
- Analyze “price of fairness”
- Algorithmic result: improved multiobjective submodular maximization
  - Improve both approximation guarantee and runtime
  - Applies to any submodular maximization problem
- Experimental analysis on real HIV prevention networks

Contributions

- Formalize measures of fairness in influence maximization
- Analyze "price of fairness"
- Algorithmic result: improved multiobjective submodular maximization
- Experimental analysis on real HIV prevention networks

Influence Maximization Background

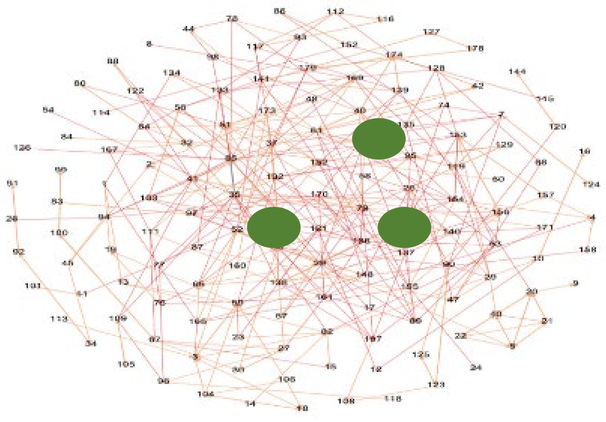

- Given:
  - Social network Graph G
  - Choose K “peer leader” nodes
- Objective:
  - Maximize expected number of influenced nodes
- *Assumption: Independent cascade model of information spread*

Defining fairness

- Notoriously difficult
- Here: look at consequences of two specific definitions

Group rationality

- Define how much influence each group should “fairly” get
- Maximize total welfare subject to those constraints

Group rationality

- What is a group’s fair share
- One notion: how much influence could they generate given a proportional share of the budget?

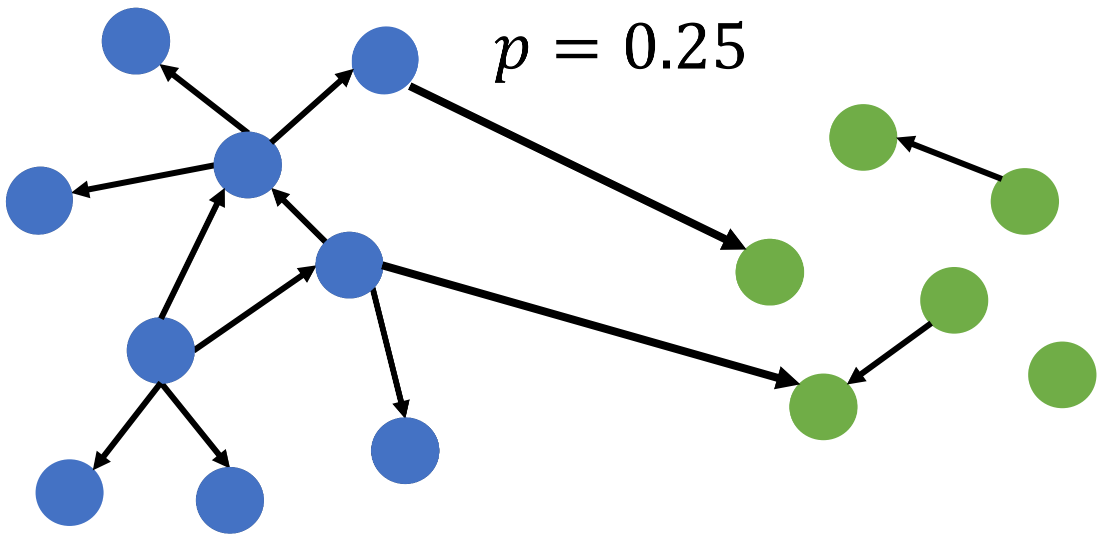

Budget: k = 5

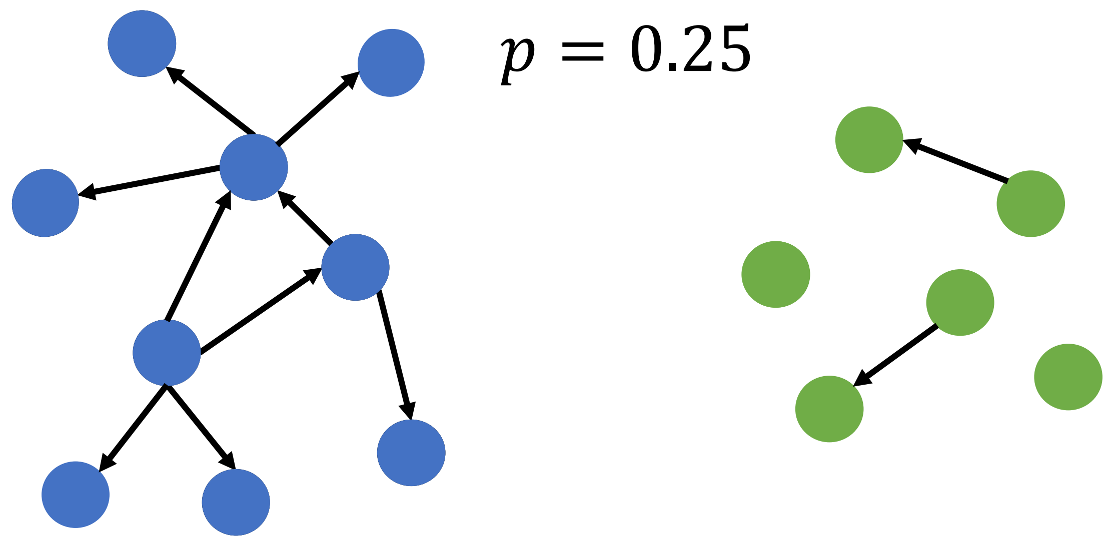

$p = 0.25$

k = 3

k = 2

$\frac{9}{15} \cdot 5$

$\frac{6}{15} \cdot 5$

Budget: k = 5

Consider each
Group independently

$p = 0.25$

Optimal value with $k = 3$:

4.5

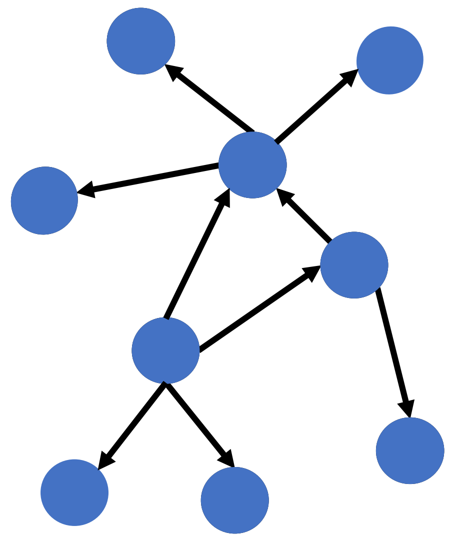

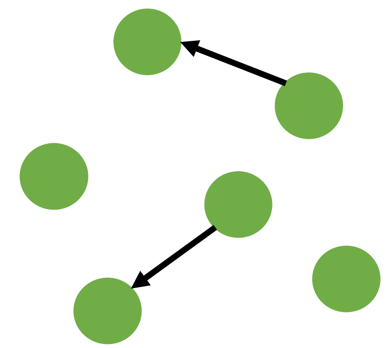

Optimal value with $k = 2$:

2.5

k = 3

k = 2

$\frac{9}{15} \cdot 5$

$\frac{6}{15} \cdot 5$

Budget: k = 5

Optimal value
with $k = 3$: 4.5

Constraint:
$f(S) \geq 4.5$

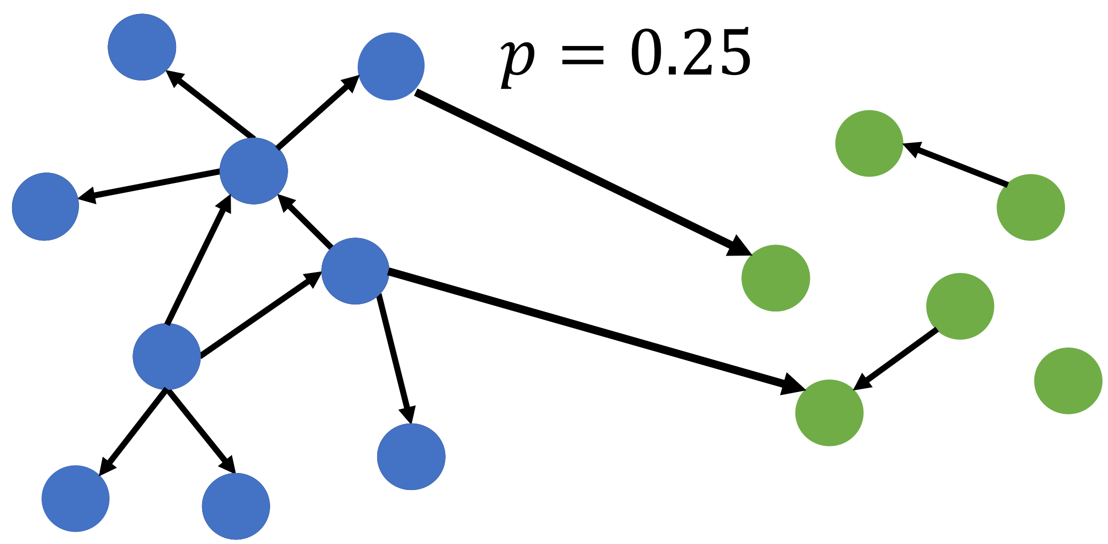

Optimal value
with $k = 2$: 2.5

Constraint:
$f(S) \geq 2.5$

Budget: k = 5

Formalization

- Input graph with groups $C_1 \dots C_m$
- Choose seed nodes S ($|S| = k$)
- Influence spreads via independent cascade model

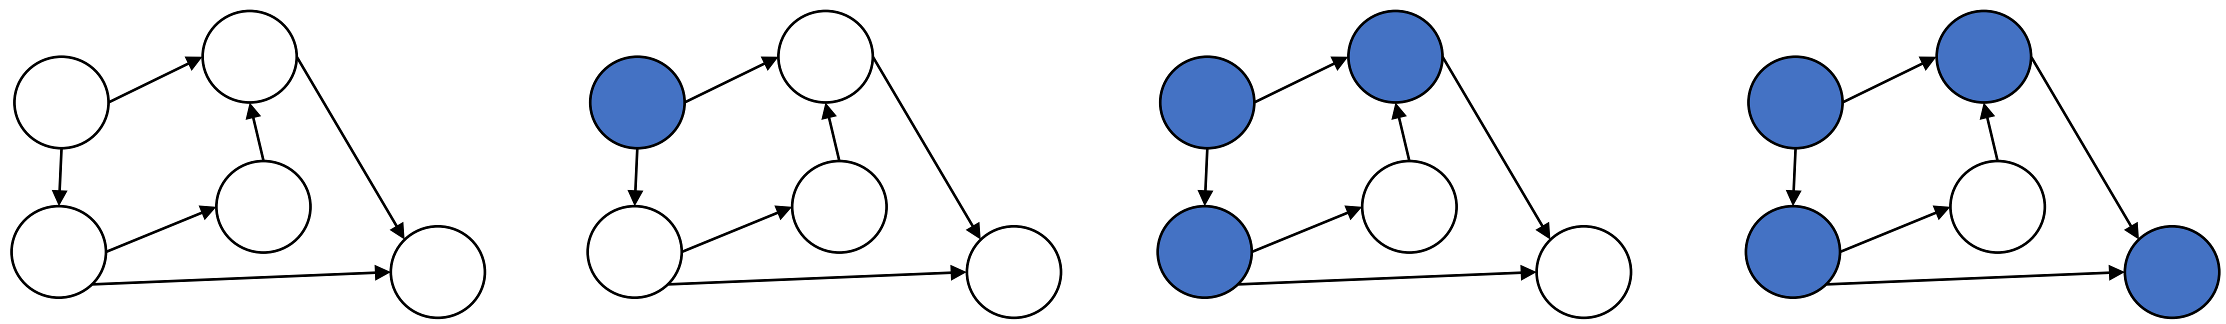

Formalization

- $f_i(S) = \mathrm{E}[\#\text{ of nodes in } \mathrm{C_i} \text{ reached}]$
- Standard objective: $\sum_{i=1}^{m} f_i$

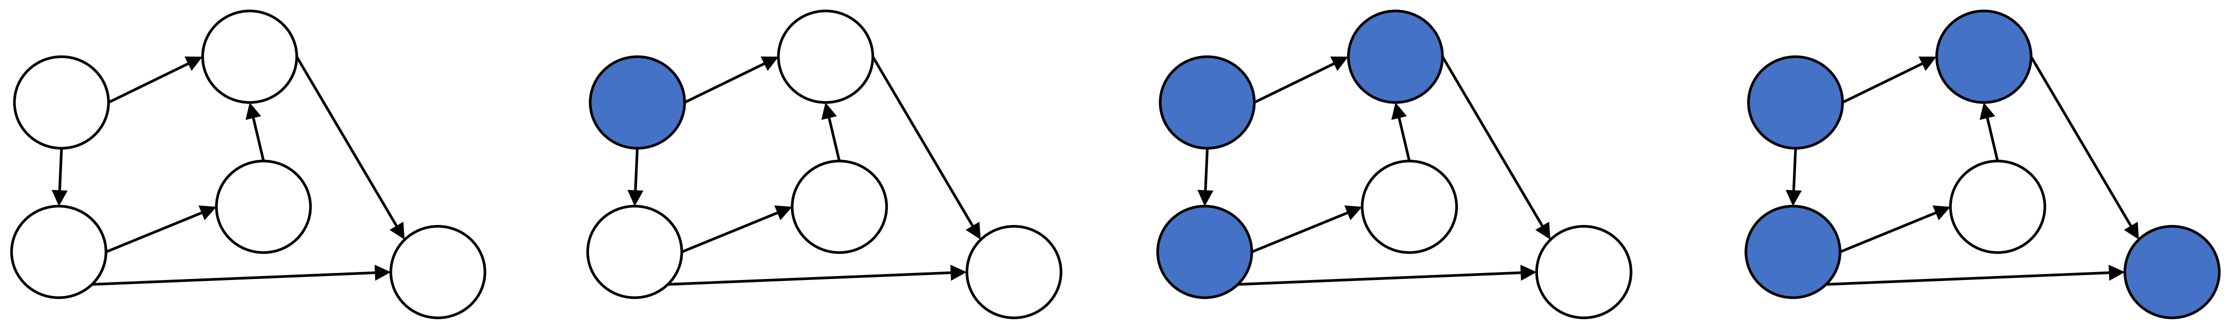

Group rationality

$$\max_S \sum_{i=1}^{m} f_i(S)$$

$$|S| \leq k$$

$$f_i(S) \geq \alpha \cdot T_i, i = 1 \ldots m$$

Group i's "fair share"

Maxmin fairness

- Maximize welfare of **least well off** group

$$\max_{|S|\leq k} \min_{i=1\ldots m} \frac{f_i(S)}{|C_i|}$$

Maxmin fairness

- Maximize welfare of **least well off** group

$$\max_{|S|\leq k} \min_{i=1\ldots m} \frac{f_i(S)}{|C_i|}$$

- Suppose 90 Blue nodes and 60 green nodes
- Suppose 45 Blue nodes influenced and 15 green nodes influenced
- Maximin fairness will require us to ensure that we improve green influence, e.g., 30 Blue and 20 green

Maxmin fairness

- Potential problem: agnostic to how well the intervention works
- Assume p=1
- K=5
- How will we ensure maxmin fairness?

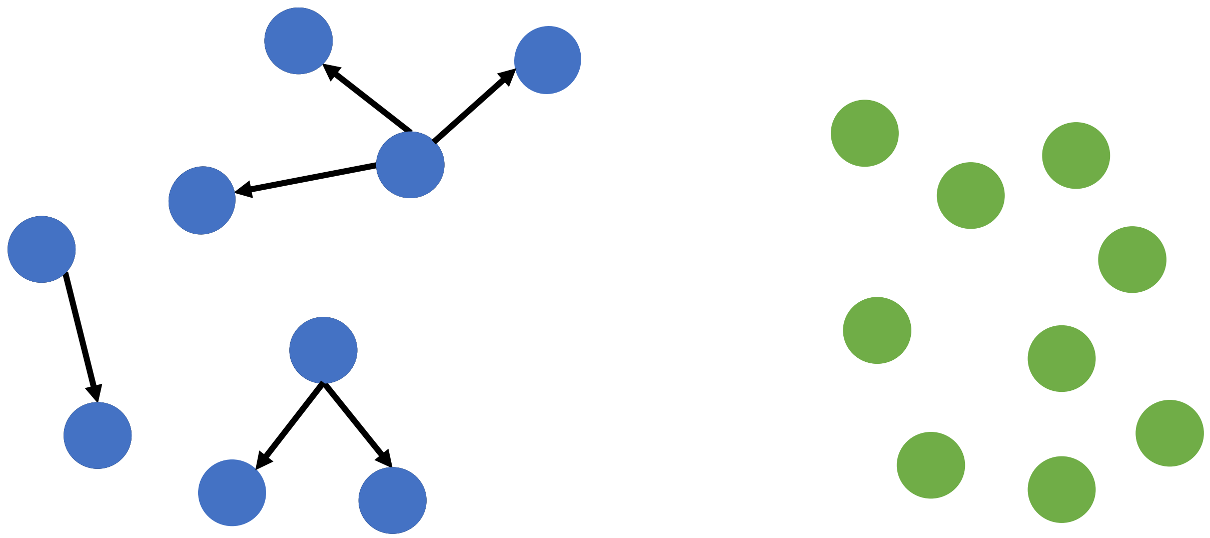

Maxmin fairness

- Potential problem: agnostic to how well the intervention works
- Assume p=1
- K=5
- How will we ensure maxmin fairness?

  Blue will get 1 seed

  Green will get 4 seeds

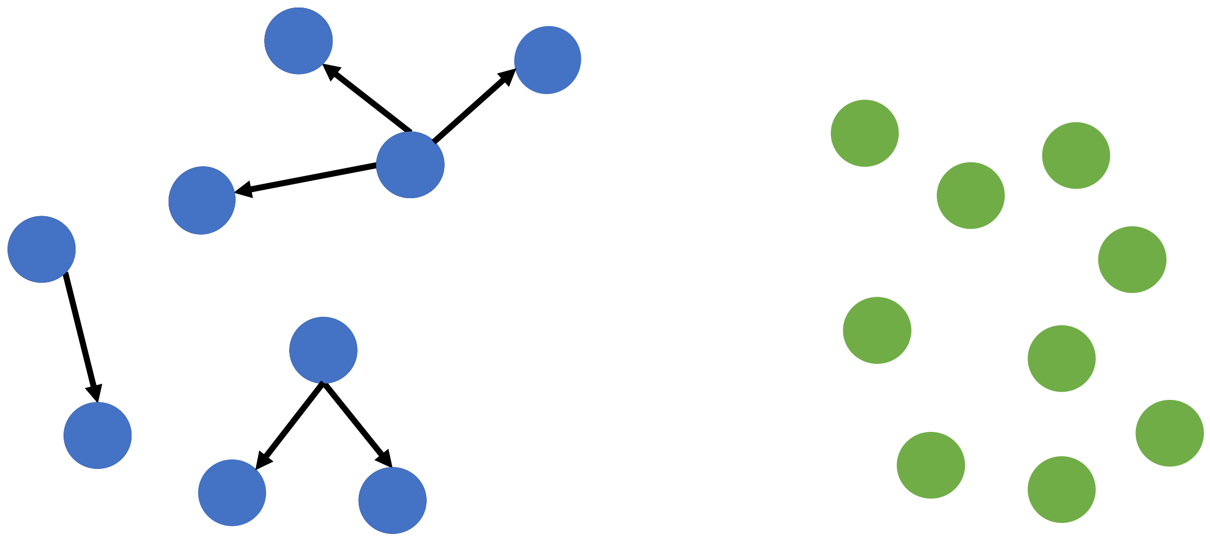

Contributions

- Formalize measures of fairness in influence maximization
- Analyze “price of fairness”
- Algorithmic result: improved multiobjective submodular maximization
- Experimental analysis on real HIV prevention networks

Price of fairness

Measures loss in total utility:

$$\frac{\sum_{i=1}^{m} f_i(S^{opt})}{\sum_{i=1}^{m} f_i(S^{fair})}$$

Price of fairness

Theorem: there exists families of graphs for which $PoF \to \infty$ as $n \to \infty$. This holds for both maxmin and group rationality.

TLDR: it's bad…

Price of fairness

Theorem: there exists families of graphs for which $PoF \to \infty$ as $n \to \infty$. This holds for both maxmin and group rationality.

TLDR: it's bad…

…But may not in practice?

Contributions

- Formalize measures of fairness in influence maximization
- Analyze "price of fairness"
- Algorithmic result: improved multiobjective submodular maximization
- Experimental analysis on real HIV prevention networks

Contributions

- Formalize measures of fairness in influence maximization
- Analyze "price of fairness"
- Algorithmic result: improved multiobjective submodular maximization
- Experimental analysis on real HIV prevention networks

Experimental results

- Four networks from HIV prevention interventions
  - ~70 nodes each
- Fairness wrt ethnicity, gender, sexual orientation, birth sex
- Compare standard greedy algorithm to both fairness concepts
  - Maxmin and Group rationality

Group Rationality:
Mean percentage violation over all groups

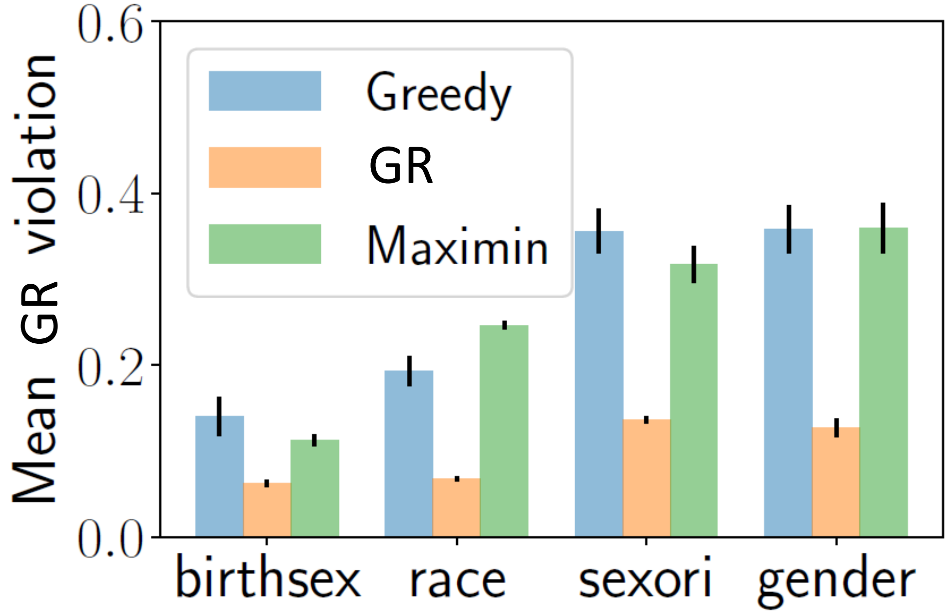

Maximin evaluation:
Min fraction influencec over all groups

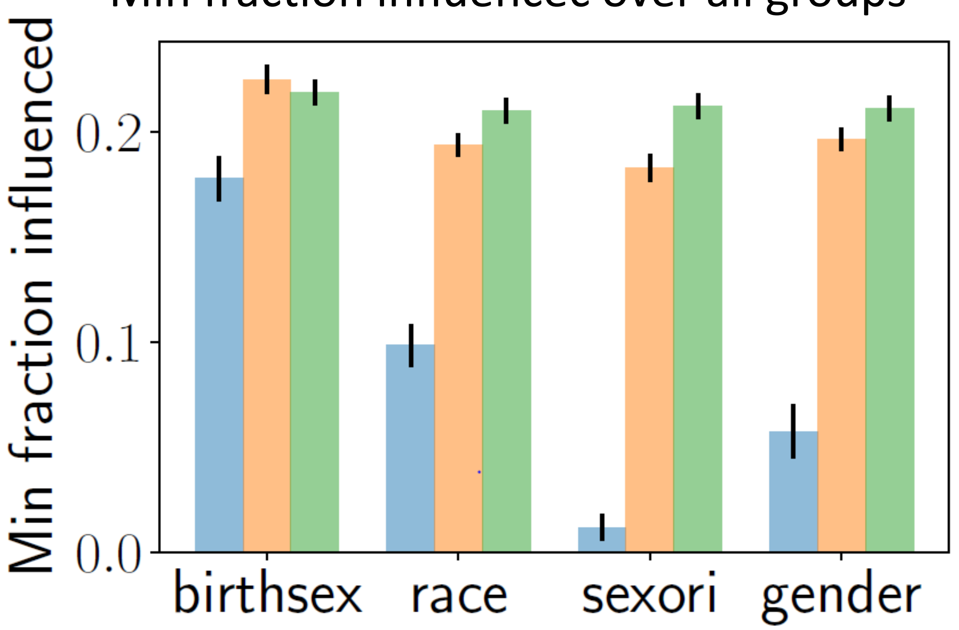

- Standard greedy algorithm sometimes neglects small groups
- Possible to substantially reduce discrepancies
- Empirically: price of fairness ≈ 1.10

Takeaways

- Influence maximization can under-allocate resources to **small groups**
- Fixing this is
  - Achievable **computationally**
  - Reasonable in terms of **overall efficiency**
- First step towards fairer interventions in socially important domains

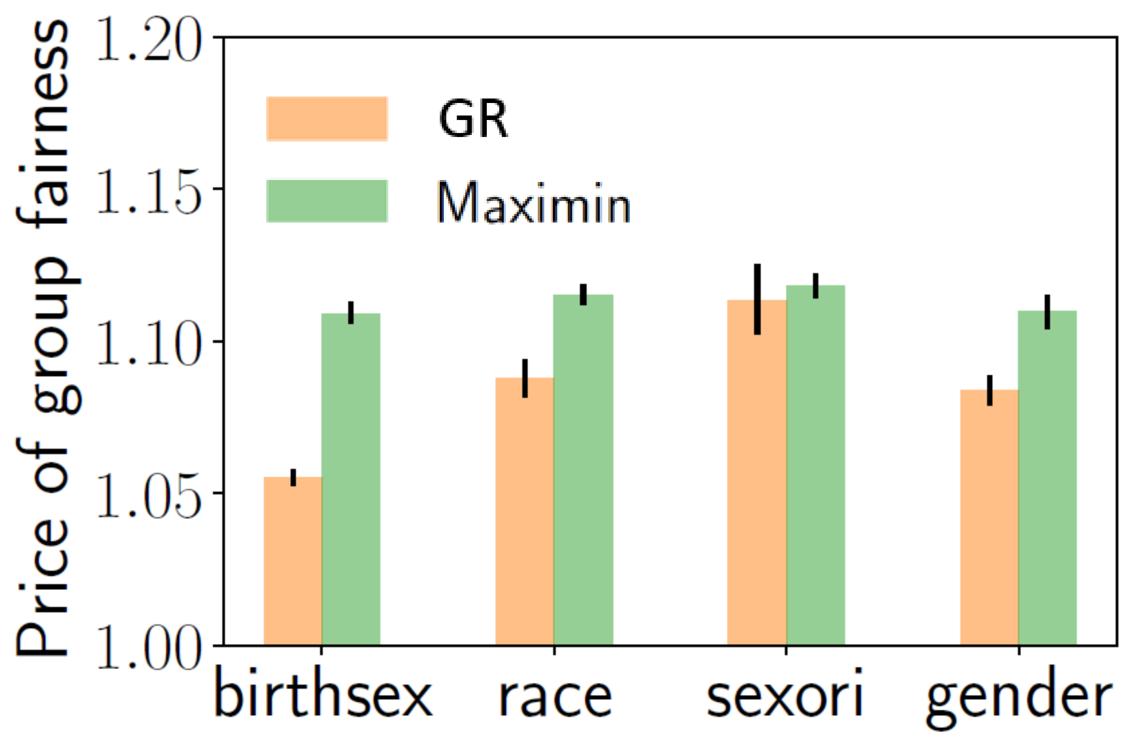

Cost of imposing fairness is relatively low (< 10% for group rationality)

Existing Greedy Algorithm

Network Name	Size	Percentage Covered by Racial Group				
		White	Black	Hisp.	Mixed	Other
SPY1	95	70	36	-	86	95
SPY2	117	78	-	42	76	67
SPY3	118	88	-	33	95	69
MFP1	165	96	77	69	73	28
MFP2	182	44	85	70	77	72

Discriminatory Coverage!

Greedy Algorithm: Tzoumas et al. Resilient monotone submodular function maximization

Numerical Results

Network Name	Size	Improvement in Minimum Percentage Covered					
		$J = 0$	$J = 1$	$J = 2$	$J = 3$	$J = 4$	$J = 5$
	95	15	16	14	10	10	8
	117	20	14	9	10	8	10
	118	20	16	16	15	11	10
	165	17	15	7	11	14	9
	182	11	12	10	9	12	12

$I = N/3$	16.6	14.6	11.2	11.0	11.0	9.8
$I = N/5$	15.0	13.8	14.0	10.0	9.0	6.7
$I = N/7$	12.2	13.4	11.4	11.4	8.2	6.4

References

- Exploring algorithmic fairness in robust graph covering problems, A. Rahmattalabi, P. Vayanos, A. Fulginiti, E. Rice, B. Wilder, A. Yadav, M. Tambe, **NeurIPS**, 2019
- K-adaptability in two-stage robust binary programming, G. Hanasusanto, D. Kuhn, W. Wiesemann. **Operations Research** 2015.
- Resilient monotone submodular function maximization, V. Tzoumas, K. Gatsis, A. Jadbabaie, G.J, Pappas, **CDC IEEE**, 2017.
- Finite adaptability in multistage linear optimization, D. Bertsimas, C. Caramanis, **IEEE Transactions on Automatic Control**, 2010.

Theoretical Results

- MMF and DC both violate population separation
- DP violates leveling down

Fairness Criteria	Fairness Notion		
	MMF	DC	DP
Population Separation	✓	✓	✗
Leveling Down	✗	✗	✓

Maxmin Fairness & Population Separation property
Group Exercise

One seed node

Optimal allocation without fairness:

Utilities of blue, red, green?

Optimal allocation with maxmin fairness?

Utilities of blue, red, green?

By imposing fairness, the utility gap should not increase!

But does it here?

DC & Population Separation property
Group Exercise

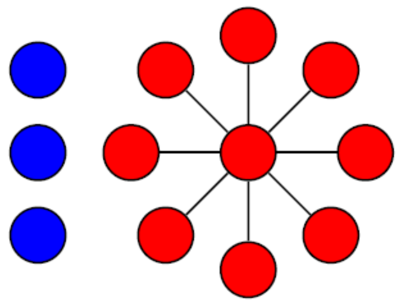

Figure 2: Companion figure to Proposition 3 of a graph with two groups: N red vertices and $N/3$ blue vertices for $N = 9$. We choose $k = 3$ and arbitrary p and $T \geq 1$. All edges are undirected, meaning that influence can spread both ways.

Three seed nodes

Optimal allocation without fairness:

Utilities of blue, red?

Optimal allocation with maxmin fairness?

Utilities of blue, red?

By imposing fairness, the utility gap should not increase!

But does it here?

Theoretical Results

- MMF and DC both suffer from population separation
- DP suffers from leveling down

Fairness Criteria	Fairness Notion		
	MMF	DC	DP
Population Separation	✓	✓	✗
Leveling Down	✗	✗	✓

Rahmattalabi

Maximin Fairness

- Incorporate the model with maximin fairness constraints
- Ensure every group receives a minimum coverage of W

$$\max_{\boldsymbol{x}\in\mathcal{X}} \min_{\boldsymbol{\xi}\in\Xi} \max_{\boldsymbol{y}\in\mathcal{Y}} \left\{ \sum_{n\in\mathcal{N}_c} y_n : y_n \leq \sum_{\nu\in\delta(n)} \xi_\nu x_n, \ \forall n \in \mathcal{N}_c \right\}$$

$$\mathcal{Y} := \left\{ \boldsymbol{y} \in \{0,1\}^N : \sum_{n\in\mathcal{N}_c} y_n \geq W|\mathcal{N}_c| \quad \forall c \in \mathcal{C}, \ \forall \boldsymbol{\xi} \in \Xi \right\}$$

- Up to 20% improvement in coverage of small groups
- Small Price of Fairness (up to 6.4% across five network samples)

Network Name	Network Size	Percentage of covered individuals by racial group				
		White	Black	Hispanic	Mixed	Other
SPY1	95	65	45	–	79	88
SPY2	117	81	–	50	72	73
SPY3	118	90	–	44	85	87
MFP1	165	85	69	42	73	64
MFP2	182	56	80	70	71	72