Fairness in Time-Critical Influence Maximization with Applications to Public Health Preventative Interventions

Aida Rahmattalabi*, Shahin jabbari†, Phebe Vayanos *, Himabindu Lakkaraju †, Milind Tambe †

Abstract

Health-promoting interventions such as gatekeeper training for suicide prevention or HIV prevention interventions have proven very successful in practice. Given the scarcity of the intervention resources in such applications, the goal is to leverage the social network structure to strategically select the individuals to train so that they can protect others or spread health information throughout the network. Indeed, the placement of an individual in a network impacts one's chances of receiving benefits from these interventions. Hence, there has been an increasing interest in the study of fairness in such problems. However, there is no clear understanding on the implications and consequences of such fairness notions or the way they can impact different sub-populations. In this work, we initiate the study of such consequences, highlight the undesirable properties of the existing notions of fairness and discuss future directions for improvements.

1 Introduction

Several health-promoting initiatives including interventions for prevention of suicide, HIV and substance abuse, propose to leverage social network structure [9, 15, 16] to strategically target individuals for interventions. In each of these applications, the goal is to select a small set of individuals who can act as peer-leaders to detect warning signs beforehand and respond appropriately (suicide prevention), or disseminate relevant information (HIV, substance abuse). Given this setting, influence maximization framework can be used to find such a set of individuals [15].

When implementing the aforementioned interventions in the real world, it is important to ensure that there is no discrimination based on protected attributes (e.g., race, ethnicity, gender, disability). Recent research has also emphasized the importance of fairness for interventions in socially sensitive settings, see e.g., [11, 15]. Several approaches have been proposed to achieve fairness in various contexts including classification [17], resource allocation [2, 4, 7], influence maximization, and graph covering [1, 13, 15]. However, there is little to no prior work on understanding the implications and consequences of each of the fairness notions proposed, or the way they can affect different sub-populations particularly in the context of influence maximization.

*University of Southern California

†Harvard University

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this work, we study the implications of several well-known notions of fairness (e.g., maximin fairness, demographic parity, and diversity constraints) in the context of influence maximization problems. Specifically, we study the vulnerabilities of the aforementioned notions and formalize why they are not readily applicable to influence maximization. To this end, we outline two properties that capture the undesirable consequences of these notions – *population separation*, and *levelling down*, and prove the aforementioned popular notions of fairness exhibit these properties in the context of influence maximization. We then propose future directions for addressing these challenges.

2 Framework

We consider the problem influence maximization problem in social networks [1]. We model a social network as a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ in which $\mathcal{N}$ is the set of all vertices (individuals) and $\mathcal{E}$ is the set of all edges (social ties). A directed edge from n to ν exists, i.e., $(n, \nu) \in \mathcal{E}$, if vertex n identifies ν as a "friend" or "social contact". We define the set of neighbors of n based on its outgoing edges.

In the influence maximization problem, a decision maker, e.g., social worker who conducts the intervention, chooses a set $\mathcal{S}$ of at most k vertices to activate. For example, in the HIV prevention, a limited set of individuals are chosen as peer-leaders to receive health-related information via a training. These individuals will then influence their peers by spreading the information through the social network. We employ the well-studied *Independent Cascade Model* [10] of influence spread. According to this model, a newly activated vertex independently activates each of its neighbors with a fixed probability $p \in (0, 1]$. This probability is typically assumed to be the same for all the edges. We further assume that influence unfolds over $T \geq 1$ time steps. At time T, the amount of influence is evaluated as the total number of vertices that are activated (have received the information). We use $\mathcal{I}(\mathcal{S})$ to denote the expected number of activated vertices with seed set $\mathcal{S} \subseteq \mathcal{N}$ after T time steps. This time-constrained influence maximization was introduced by Chen, Lu, and Zhang [3]. We adopt this model as in most applications of interest the spread of influence is time-critical [13]. In addition, this model agrees with the observation that the spread of influence in social networks slows down [8].

Formally, we define the influence maximization problem as follows:

$$\max_{\mathcal{S}\subseteq\mathcal{N},\,|\mathcal{S}|\leq k} \quad \mathcal{I}(\mathcal{S}). \tag{1}$$

This is a general model that encompasses as a special case the classic influence maximization problem ($T = |\mathcal{N}|$) and the graph covering problem ($T = 1, p = 1$).

Fairness Considerations When maximizing $\mathcal{I}$, (near) optimal interventions end up mirroring the differences in degree of connectedness of different groups and can result in discriminatory outcomes [13]. However, it is important to ensure that there is no discrimination based on protected attributes. For example, in the HIV prevention domain one needs to ensure that racial minorities or individuals with LGBTQ identity are not disproportionately excluded. Discrimination with respect to protected attributes and socially salient groups has recently been studied in the context of (robust) covering problems [13] and for the (stochastic) influence maximization [1, 15].

Mathematically, each vertex (individual) $n \in \mathcal{N}$ is characterized by a set of attributes (protected characteristics) such as race or gender, for which fair treatment is important. Based on these protected characteristics, one can partition $\mathcal{N}$ into $C \geq 2$ disjoint groups $\mathcal{N}_c$, $c \in \mathcal{C} := \{1,\ldots,C\}$, such that $\mathcal{N}_1 \cup \cdots \cup \mathcal{N}_C = \mathcal{N}$. We use $\mathcal{I}_c(\mathcal{S})$ to denote the expected number of influenced vertices in group c under the seed set $\mathcal{S}$ and we define each group's utility as the influence normalized by the size of the group under seed set $\mathcal{S}$. Precisely, for any $\mathcal{S}$, we can define the utility vector $\boldsymbol{u}(\mathcal{S}) \in [0,1]^C$ as $\boldsymbol{u}(\mathcal{S}) := (\mathcal{I}_1(\mathcal{S})/|\mathcal{N}_1|,\ldots,\mathcal{I}_C(\mathcal{S})/|\mathcal{N}_C|)$, where the cth component $\boldsymbol{u}_c$ indicates the utility of group c. Furthermore, we use $\mathcal{U}$ to denote the set of all possible utilities that the groups can achieve simultaneously for a fixed budget k i.e., $\mathcal{U}(k) := \cup_{|\mathcal{S}|\leq k}\ \boldsymbol{u}(\mathcal{S})$. When clear from the context, we drop the dependence on $\mathcal{S}$ and k for convenience.

As a result, we can write Problem (1) equivalently as

$$\max_{\boldsymbol{u}} \quad \sum_{c\in\mathcal{C}} |\mathcal{N}_c|\boldsymbol{u}_c :\ \boldsymbol{u} \in \mathcal{U}. \tag{2}$$

We use $\mathcal{U}^{\text{opt}}$ to denote the set of optimal solutions of Problem (2) and we refer to $\boldsymbol{u} \in \mathcal{U}^{\text{opt}}$ as an optimal utility.

Problem (2) does not have any fairness considerations. A common approach to incorporate fairness is to impose constraints on the utility vectors $\boldsymbol{u}$. We use $\mathcal{F}$ to collect the fairness constraints and we define the problem

$$\max_{\boldsymbol{u}} \quad \sum_{c\in\mathcal{C}} |\mathcal{N}_c|\boldsymbol{u}_c :\ \boldsymbol{u} \in \mathcal{U},\ \boldsymbol{u} \in \mathcal{F}. \tag{3}$$

We use $\mathcal{U}^{\text{fair}}$ to denote the set of optimal solutions to Problem (3) and we refer to $\boldsymbol{u} \in \mathcal{U}^{\text{fair}}$ as an optimal fair utility. We note that depending on the fairness notion, the set F can be precisely defined. In the following section, we discuss several notions of fairness with their corresponding sets $\mathcal{F}$.

3 Notions of Fairness

We now study fairness in the influence maximization problem. We start by reviewing three notions of fairness. Henceforth, we will refer to the group with the minimum (resp. maximum) utility as the worst-off (resp. best-off) group.

Maximin Fairness Maximin fairness is a generalization of Rawlsian theory of justice [14], according to which the minimum utility that the groups obtain should be maximized. Following the approach in [15], we can define the set $\mathcal{F}$ by requiring a minimum level of utility for every 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.*

Demographic Parity Demographic parity requires that independently of their group membership, the utility of each group be the same. This notion formalizes the legal doctrine of disparate impact [17]. A generalization of this notion has been introduced by Hardt et al. [6] in the classification context as equality of opportunity. Their definition caters for the case that some individuals are better suited for an intervention. We note that in the applications of interest, every individual/group is equally qualified to receive the benefit from the intervention, e.g., everyone should be protected by a gatekeeper or should receive health-promoting information. In effect, in our resource allocation setting demographic parity and equality of opportunity are equivalent.

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

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

In Definition 2, δ is a parameter that determines the strictness in fairness where smaller δ values impose more strict fairness constraints. We note that for $\delta \geq 1$ the constraints $\mathcal{F}$ in Definition 2 hold trivially, thus, we require $\delta \in [0,1)$.

MMF and DP do not depend on the structure of the network $\mathcal{G}$. Recently, Tsang et al. [15] proposed a notion of fairness called diversity constraint that explicitly takes the network structure into account.

Diversity Constraints Diversity constraints require that no group be better-off (obtain a higher utility) with an allocation of resources proportional to their size and re-allocating them internally. Diversity constraints are defined as follows.

Definition 3 (Diversity Constraints [15])**.** *Let $k_c = \lceil k|\mathcal{N}_c|/N\rceil$ and $\mathcal{U}_c(k) := \cup_{|\mathcal{S}|\leq k, \mathcal{S}\subseteq\mathcal{N}_c}\ \boldsymbol{u}(\mathcal{S})\ \forall c \in \mathcal{C}$ be the set of utility vectors when the seed vertices are chosen from group c. The utility of each group must be at least $\overline{\boldsymbol{u}}_c = \max_{\boldsymbol{u}\in\mathcal{U}_c(k_c)} \boldsymbol{u}_c\ \forall c \in \mathcal{C}$ to satisfy the* diversity constraints (DC). *Formally, the set $\mathcal{F}$ for* DC *is defined as*

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

4 Fairness Criteria

We now discuss implications of incorporating the notions of fairness introduced in section 3. One concern may be that imposing fairness can increase the utility gap between groups, where for any utility vector $\boldsymbol{v} \in \mathcal{U}$, utility gap is defined as $\Delta\boldsymbol{v} := \max_{c,c'\in\mathcal{C}}(\boldsymbol{v}_c - \boldsymbol{v}_{c'})$. This situation is undesirable as it causes a greater separation in the utilities of the groups. Fish et al. [5] studied a similar notion as rich get richer but their definition is restricted to only two groups.

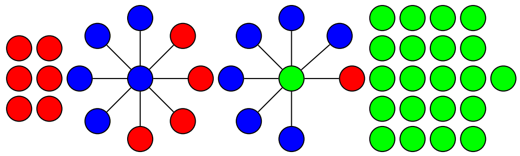

Figure 1: Companion figure to Proposition 1 for the case of $p = 1$. The network consists of three groups: red, blue and green. The edges are undirected so the influence can spread both ways. For arbitrary p, the number of isolated green vertices should scale to $\lceil 21/p \rceil$.

Definition 4 (Population Separation). *A fairness notion exhibits population separation if there exists a graph G, time horizon T, budget k and propagation probability p such that*

$$\exists \boldsymbol{u} \in \mathcal{U}^{\text{opt}}, \boldsymbol{w} \in \mathcal{U}^{\text{fair}} : \Delta \boldsymbol{u} < \Delta \boldsymbol{w}.$$

Remark 1. *The population separation is particularly undesirable when not only do the best-off and worst-off groups not change in the fair solution, but also utility gap widens.*

Our next definition is motivated by the desire that fairness should not result in degradation of the utilities of the groups. This idea appears as levelling down objection in political philosophy [12]. We introduce two variants of this property.

Definition 5 (Weak Levelling Down). *A fairness notion exhibits* weak levelling down *if there exists a graph G, time horizon T, budget k and propagation probability p such that*

$$\exists \boldsymbol{u} \in \mathcal{U}^{\text{opt}}, \boldsymbol{w} \in \mathcal{U}^{\text{fair}} : \boldsymbol{w}_c \leq \boldsymbol{u}_c \;\; \forall c \in \mathcal{C} \text{ and } \; \exists c \; \boldsymbol{w}_c < \boldsymbol{u}_c.$$

Definition 5 is related to the Pareto optimality axiom studied by Bertsimas, Farias, and Trichakis [2] which ensures that imposing fairness does not result in wastage of utilities. In a stronger notion, all utilities simultaneously degrade.

Definition 6 (Strong Levelling Down). *A fairness notion exhibits* strong levelling down *if there exists a graph G, time horizon T, budget k and propagation probability p such that*

$$\exists \boldsymbol{u} \in \mathcal{U}^{\text{opt}}, \boldsymbol{w} \in \mathcal{U}^{\text{fair}} : \boldsymbol{w}_c < \boldsymbol{u}_c \;\; \forall c \in \mathcal{C}.$$

Proposition 1. *For $C > 2$,* MMF *exhibits the population separation property.*

Proof. First, we note that if $C = 2$, MMF does not exhibit the population separation property. This is because the utility of any of the worst-off groups in the optimal solution does not decrease after enforcing fairness. In this case, from population separation, it follows that the utility of the best-off group should increase after enforcing fairness which is not possible as the total influence of the optimal fair solution cannot be strictly higher than the total influence for the optimal solution. For $C > 2$, we prove the statement via the example in Figure 1 which depicts a network with three groups: blue, green and red. We fix $k = 1$ and require $T \geq 1$ and $p > 3/4$. The graph corresponds to the case where $p = 1$ but the example will hold for arbitrary p by setting the number of isolated green vertices to be $\lceil 21/p \rceil$. The optimal solution of Problem (2) targets the center of the bigger star component. Thus, the utilities of blue, green and red will be $(1 + 4p)/11$, $4p/11$ and 0, respectively. This results in utility gap equal to $(1+4p)/11$. By imposing MMF, the optimal fair solution selects the center of the smaller star component and the optimal fair utilities of blue, green and red will be $6p/11$, $p/11$ and $1/(1 + \lceil 21/p \rceil)$, respectively and we observe a utility gap $6p/11 - 1/(1 + \lceil 21/p \rceil) \geq 6p/11 - 1/22 > (1 + 4p)/11$, where we used $p > 3/4$ in the last inequality. ■

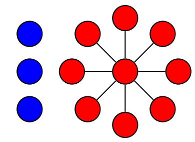

Figure 2: Companion figure to Proposition 3 of a graph with two groups: N red vertices and $N/3$ blue vertices for $N = 9$. We choose $k = 3$ and arbitrary p and $T \geq 1$. All edges are undirected, meaning that influence can spread both ways.

Proposition 2. *For all $\delta \in [0, 1)$,* DP *does not exhibit population separation.*

Proof. We prove by contradiction. We consider the problem instances that $\exists \boldsymbol{u} \in \mathcal{U}^{\text{opt}} : \Delta \boldsymbol{u} > \delta$ (otherwise there is no need to impose DP constraints as all the optimal solutions satisfy DP constraints). Suppose after imposing fairness constraints we observe population separation, i.e., $\exists \boldsymbol{w} \in \mathcal{U}^{\text{fair}} : \Delta \boldsymbol{u} < \Delta \boldsymbol{w}$. According to the definition of DP, it follows that $\Delta \boldsymbol{w} \leq \delta$. It must also hold that $\Delta \boldsymbol{u} < \delta$. This, however, contradicts the assumption that $\Delta \boldsymbol{u} > \delta$. ■

Proposition 3. DC *exhibits population separation.*

Proof. Consider the graph $\mathcal{G}$ as in Figure 2 consisting of two groups blue and red with size $N/3$ and N, respectively. Suppose $k = 3$ and p and $T \geq 1$ are arbitrary. Without DC, an optimal solution places one seed vertex at the center of the red group and allocates the remaining 2 vertices to the blue group. Thus, the utility of the red and blue groups will be equal to $(1 + (N-1)p)/N$ and $6/N$. Since there is no edge between the red and blue groups, DC (See Definition 3) reduces to how to optimally choose one seed vertex from blue and the remaining two from the red group. After imposing DC, the utility of the red and blue groups will be equal to $(2 + (N-2)p)/N$ and $3/N$. Therefore, the utility gap between the two groups increases. ■

Proposition 2 indicates that DP exhibits behavior opposite to MMF or DC in terms of population separation. A natural question is that whether imposing both sets of constraints can enhance MMF or DC. Similar to the proof of Proposition 2, one can prove that MMF/DC combined with DP does not exhibit population separation. Next, we investigate the strong levelling down property, where we show MMF and DC do not exhibit strong levelling down, whereas DP suffers from the strong (and hence the weak) levelling down.

Proposition 4. *Consider a general fairness notion as a set of constraints $\mathcal{F} = \{\boldsymbol{u} \in [0,1]^C : \boldsymbol{u}_c \geq v_c \;\; \forall c \in \mathcal{C}\}$ where $v_c \; \forall c \in \mathcal{C}$ are arbitrary lower-bound values. The considered fairness notion does not exhibit strong levelling down.*

Proof. Suppose the contrary holds, meaning that $\exists \boldsymbol{u} \in \mathcal{U}^{\text{opt}}, \boldsymbol{w} \in \mathcal{U}^{\text{fair}} : \boldsymbol{w}_c < \boldsymbol{u}_c \;\; \forall c \in \mathcal{C}$ We consider two cases. First, assume $\exists c \in \mathcal{C} : \boldsymbol{u}_c < v_c$. By $\boldsymbol{w}_c \geq v_c$, it must be that $\boldsymbol{u}_c < \boldsymbol{w}_c$ which is a contradiction. Second, $\forall c \in \mathcal{C} : \boldsymbol{u}_c \geq v_c$. This, however, contradicts the optimality of $\boldsymbol{w}$ as there exists a solution that satisfies the fairness constraints and has a strictly better objective value. ■

Corollary 1. MMF *and* DC *do not exhibit the strong levelling down property.*

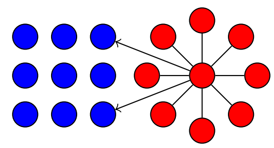

Figure 3: Companion figure to Proposition 5 where we choose $\max(3p/(p-\delta),(p-2p^2)/(\delta-p^2)) < N$, $\delta < p < \min(\sqrt{\delta}, 0.5)$, $k=2$ and $T>1$. The network consists of two groups red and blue, each of size N. All edges except the two shown by arrows are undirected meaning that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

Proposition 5. *For all* $\delta \in [0,1)$*,* DP *exhibits the strong levelling down property.*

Proof. Consider a graph $\mathcal{G}$ as shown in Figure 3 consisting of two groups, blue and red, each of size N. We choose an arbitrary δ, to reflect the arbitrary strictness of a decision maker. Let $\delta < p < \min(\sqrt{\delta}, 0.5)$, $k=2$ and $T>1$ and N be large enough, i.e., $\max(3p/(p-\delta),(p-2p^2)/(\delta-p^2)) < N$. Without DP constraints, an optimal solution chooses the center of the star and one of the blue vertices. The utility of red and blue will be $(1+(N-1)p)/N$ and $(1+2p)/N$, respectively. In this case, the utility gap exceeds δ. By imposing DP, an optimal solution is to choose one vertex from the periphery of the red group and one vertex from the isolated blue vertices. The utilities of red and blue will be $\left(1+p+p^2(N-2)\right)/N$ and $1/N$, respectively. Given the range of N, the utility gap is less than δ yet the utility is strictly smaller than the case without DP constraints. ■

These results raise the question of whether combining different notions of fairness can help avoid the pitfalls of each definition. This, however, is not possible. For instance, consider combining MMF and DP. Using the same example as in the proof of Proposition 5, we observe that the choice of seed vertices after imposing DP is in fact optimal if MMF constraints are added ($W = 1/N$), proving that DP combined with MMF still suffers from levelling down. Similar issues are observed for the other combinations of the fairness notions. We leave the details out for space considerations.

5 Discussion

In this work, we studied the implications of popular notions of fairness in the context of influence maximization problems. Our research demonstrates that these notions when applied in the context of influence maximization result in undesirable consequences such as a) *population separation*, where the utility gap between the best-off and worst-off groups increases, and b) *levelling down* which corresponds to the sub-optimal allocation of resources, which in turn results in degradation of the utilities of certain groups (which could have been avoided). While such undesirable consequences might not occur in a classification setting, our research shows that they pose challenges in influence maximization problems. One of the reasons for this difference in behavior is the fact that popular notions of fairness do not account for graph structure. Therefore, it is important to come up with novel notions of fairness which incorporate metrics pertaining to the graph structure. This opens up interesting directions for future work.

References

[1] Ali, J.; Babaei, M.; Chakraborty, A.; Mirzasoleiman, B.; Gummadi, K.; and Singla, A. 2019. On the fairness of time-critical influence maximization in social networks. *arXiv preprint:1905.06618*.

[2] Bertsimas, D.; Farias, V.; and Trichakis, N. 2011. The price of fairness. *Operations Research* 59(1):17–31.

[3] Chen, W.; Lu, W.; and Zhang, N. 2012. Time-critical influence maximization in social networks with time-delayed diffusion process. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 592–598. AAAI Press.

[4] Conitzer, V.; Freeman, R.; Shah, N.; and Vaughan, J. 2019. Group fairness for the allocation of indivisible goods. *Proceedings of the AAAI Conference on Artificial Intelligence* 33:1853–1860.

[5] Fish, B.; Bashardoust, A.; Boyd, D.; Friedler, S.; Scheidegger, C.; and Venkatasubramanian, S. 2019. Gaps in information access in social networks? In *Proceedings of the 28th World Wide Web Conference*, 480–490.

[6] Hardt, M.; Price, E.; Srebro, N.; et al. 2016. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, 3315–3323.

[7] Iancu, D. A., and Trichakis, N. 2014. Fairness and efficiency in multiportfolio optimization. *Operations Research* 62(6):1285–1301.

[8] Iribarren, J. L., and Moro, E. 2009. Impact of human activity patterns on the dynamics of information diffusion. *Physical Review Letters* 103(3):038702.

[9] Isaac, M.; Elias, B.; Katz, L.; Belik, S.-L.; Deane, F.; Enns, M.; Sareen, J.; and members), S. C. S. P. T. . 2009. Gatekeeper training as a preventative intervention for suicide: a systematic review. *The Canadian Journal of Psychiatry* 54(4):260–268.

[10] Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 137–146.

[11] Kube, A.; Das, S.; and Fowler, P. J. 2019. Allocating interventions based on predicted outcomes: A case study on homelessness services. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference 2019*, 622–629.

[12] Parfit, D. 1997. Equality and priority. *Ratio* 10(3):202–221.

[13] Rahmattalabi, A.; Vayanos, P.; Fulginiti, A.; Rice, E.; Wilder, B.; Yadav, A.; and Tambe, M. 2019. Exploring algorithmic fairness in robust graph covering problems. In *Advances in Neural Information Processing Systems*.

[14] Rawls, J. 2009. *A theory of justice*. Harvard university press.

[15] Tsang, A.; Wilder, B.; Rice, E.; Tambe, M.; and Zick, Y. 2019. Group-fairness in influence maximization. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 5997–6005.

[16] Valente, T.; Ritt-Olson, A.; Stacy, A.; Unger, J.; Okamoto, J.; and Sussman, S. 2007. Peer acceleration: effects of a social network tailored substance abuse prevention program among high-risk adolescents. *Addiction* 102(11):1804–1815.

[17] Zafar, M. B.; Valera, I.; Gomez Rodriguez, M.; and Gummadi, K. P. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, 1171–1180. International World Wide Web Conferences Steering Committee.