

Fair Influence Maximization: A Welfare Optimization Approach

Aida Rahmattalabi,¹ Shahin Jabbari,² Himabindu Lakkaraju,² Phebe Vayanos,¹ Max Izenberg,³
Ryan Brown,³ Eric Rice,¹ Milind Tambe²

¹University of Southern California

²Harvard University

³RAND Corporation

{rahmatta, phebe.vayanos, eric}@usc.edu, {jabbari, milind_tambe}@seas.harvard.edu, hlakkaraju@hbs.edu, {izenberg, rbrown}@rand.org

Abstract

Several behavioral, social, and public health interventions, such as suicide/HIV prevention or community preparedness against natural disasters, leverage social network information to maximize outreach. Algorithmic influence maximization techniques have been proposed to aid with the choice of “peer leaders” or “influencers” in such interventions. Yet, traditional algorithms for influence maximization have not been designed with these interventions in mind. As a result, they may disproportionately exclude minority communities from the benefits of the intervention. This has motivated research on *fair influence maximization*. Existing techniques come with two major drawbacks. First, they require committing to a single fairness measure. Second, these measures are typically imposed as strict constraints leading to undesirable properties such as wastage of resources. To address these shortcomings, we provide a principled characterization of the properties that a fair influence maximization algorithm should satisfy. In particular, we propose a framework based on social welfare theory, wherein the cardinal utilities derived by each community are aggregated using the isoelastic social welfare functions. Under this framework, the trade-off between fairness and efficiency can be controlled by a single *inequality aversion* design parameter. We then show under what circumstances our proposed principles can be satisfied by a welfare function. The resulting optimization problem is monotone and submodular and can be solved efficiently with optimality guarantees. Our framework encompasses as special cases leximin and proportional fairness. Extensive experiments on synthetic and real world datasets including a case study on landslide risk management demonstrate the efficacy of the proposed framework^{1,2}

Introduction

The success of many behavioral, social, and public health interventions relies heavily on effectively leveraging social networks (Isaac et al. 2009; Valente et al. 2007; Tsang et al. 2019). For instance, health interventions such as suicide/HIV prevention (Yonemoto et al. 2019) and community preparedness against natural disasters involve finding a small set of

well-connected individuals who can act as peer-leaders to detect warning signals (suicide prevention) or disseminate relevant information (HIV prevention or landslide risk management). The influence maximization framework has been employed to find such individuals (Wilder et al. 2020). However, such interventions may lead to discriminatory solutions as individuals from racial minorities or LGBTQ communities may be disproportionately excluded from the benefits of the intervention (Rahmattalabi et al. 2019; Tsang et al. 2019).

Recent work has incorporated fairness directly into influence maximization by proposing various notions of fairness such as maximin fairness (Rahmattalabi et al. 2019) and diversity constraints (Tsang et al. 2019). Maximin fairness aims at improving the minimum amount of influence that any community receives. On the other hand, diversity constraints, inspired by the game theory literature, ensure that each community is at least as well-off had they received their share of resources proportional to their size and allocated them internally. Each of these notions offers a unique perspective on fairness. However, they also come with drawbacks. For example maximin fairness can result in significant degradation in total influence due to its stringent requirement to help the worst-off group as much as possible, where in reality it may be hard to spread the influence to some communities due to their sparse connections. On the other hand, while the diversity constraints aim at taking the community’s ability in spreading influence into account, it does not explicitly account for reducing inequality (i.e., does not exhibit *inequality aversion*). Consequently, there is no universal agreement on what fairness means and in fact, it is widely known that fairness is domain dependent (Narayanan 2018). For example, excluding vulnerable communities from suicide prevention might have higher negative consequences compared to interventions promoting a healthier lifestyle.

Building on cardinal social welfare theory from the economics literature and principles of social welfare, we propose a principled characterization of the properties of social influence maximization solutions. In particular, we propose a framework for fair influence maximization based on social welfare theory, wherein the cardinal utilities derived by each community are aggregated using the isoelastic social welfare functions (Bergson 1938). Isoelastic functions are in the gen-

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Full version at <https://arxiv.org/abs/2006.07906>

²The first two authors contributed equally.

eral form of u^α/α , $\alpha < 1$, $\alpha \neq 0$ and $\log u$, $\alpha = 0$ where α is a constant and controls the aversion to inequality and u is the utility value. They are used to measure the *goodness* or *desirability* of a utility distribution. However, due to the structural dependencies induced by the underlying social network, i.e., between-community and within-community edges, social welfare principles cannot be directly applied to our problem. Our contributions are as follows: (i) We extend the cardinal social welfare principles including the *transfer principle* to the influence maximization framework, which is otherwise not applicable. We also propose a new principle which we call *utility gap reduction*. This principle aims to avoid situations where high aversion to inequality leads to even more utility gap, caused by between-community influence spread; (ii) We generalize the theory regarding these principles and show that for all problem instances, there does not exist a welfare function that satisfies all principles. Nevertheless, we show that if all communities are disconnected from one another (no between-community edges), isoelastic welfare functions satisfy all principles. This result highlights the importance of network structure, specifically between-community edges; (iii) Under this framework, the trade-off between fairness and efficiency can be controlled by a single *inequality aversion* parameter α . This allows a decision-maker to effectively trade-off quantities like utility gap and total influence by varying this parameter in the welfare function. We then incorporate these welfare functions as objective into an optimization problem to rule out undesirable solutions. We show that the resulting optimization problem is monotone and submodular and, hence, can be solved with a greedy algorithm with optimality guarantees; (iv) Finally, we carry out detailed experimental analysis on synthetic and real social networks to study the trade-off between total influence spread and utility gap. In particular, we conduct a case study on the social network-based landslide risk management in Sitka, Alaska. We show that by choosing α appropriately we can flexibly control utility gap (4%-26%) and the resulting influence degradation (36% - 5%).

Related Work. Recent work has incorporated fairness directly into the influence maximization framework by relying on either Rawlsian theory (Rawls 2009; Rahmattalabi et al. 2019), game theoretic principles (Tsang et al. 2019) or equality based notions (Stoica, Han, and Chaintreau 2020). Equality based approaches strive for equal outcome across different communities. In general, strict equality is hard to achieve and may lead to wastage of resources. This is amplified in influence maximization as different communities have different capacities in being influenced (e.g., marginalized communities are hard to reach). We discuss the other two approaches in more details in Sections –. Social welfare functions have been used within the economic literature to study trade-offs between equality and efficiency (Sen 1997) and have been widely adopted in different decision making areas including health (Abasolo and Tsuchiya 2004). Recently, Heidari et al. (2018) proposed to study similar ideas for regression problems. The classical social welfare theory, however, does not readily extend to our setting due to dependencies induced by the between-community connections. Extending those principles is a contribution of our work. See the full version for a

thorough review of related work.

Problem Formulation

We use $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ to denote a graph (or network) in which \mathcal{V} is the set of N vertices and \mathcal{E} is the set of edges. In the *influence maximization problem*, a decision-maker chooses a set of at most K vertices to influence (or activate). The selected vertices then spread the influence in rounds according to the *Independent Cascade Model* (Kempe, Kleinberg, and Tardos 2003).³ Under this model, each newly activated vertex spreads the influence to its neighbors independently and with a fixed probability $p \in [0, 1]$. The process continues until no new vertices are influenced. We use \mathcal{A} to denote the initial set of vertices, also referred to as influencer vertices. The goal of the decision-maker is to select a set \mathcal{A} to maximize the expected number of vertices that are influenced at the end of this process. Each vertex of the graph belongs to one of the disjoint communities (empty intersection) $c \in \mathcal{C} := \{1, \dots, C\}$ such that $\mathcal{V}_1 \cup \dots \cup \mathcal{V}_C = \mathcal{V}$ where \mathcal{V}_c denotes the set of vertices that belong to community c . This partitioning can be induced by, e.g., the intersection of a set of (protected) attributes such as race or gender for which fair treatment is important. We use N_c to denote the size of community c , i.e., $N_c = |\mathcal{V}_c|$. Furthermore, communities may be disconnected, in which case $\forall c, c' \in \mathcal{C}$ and $\forall v \in \mathcal{V}_c, v' \in \mathcal{V}_{c'}$, there is no edge between v and v' (i.e., $(v, v') \notin \mathcal{E}$). We define $\mathcal{A}^* := \{\mathcal{A} \subseteq \mathcal{V} \mid |\mathcal{A}| \leq K\}$ as the set of budget-feasible influencers. Finally, for any choice of influencers $\mathcal{A} \in \mathcal{A}^*$, we let $u_c(\mathcal{A})$ denote the utility, i.e., the expected fraction of the influenced vertices of community c , where the expectation is taken over randomness in the spread of influence. The standard influence maximization problem solves the optimization problem

$$\underset{\mathcal{A} \in \mathcal{A}^*}{\text{maximize}} \sum_{c \in \mathcal{C}} N_c u_c(\mathcal{A}). \quad (1)$$

When clear from the context, we will drop the dependency of $u_c(\mathcal{A})$ on \mathcal{A} to minimize notational overhead.

Existing Notions of Fairness

Problem (1) solely attempts to maximize the total influence which is also known as the *utilitarian* approach. Existing fair influence maximization problems are variants of Problem (1) involving additional constraints. We detail these below. See the full version for more discussion.

Maximin Fairness (MMF). Based on the Rawlsian theory (Rawls 2009), MMF (Tsang et al. 2019) aims to maximize the utility of the worst-off community. Precisely, MMF only allows $\mathcal{A} \in \mathcal{A}^*$ that satisfy the following constraint

$$\min_{c \in \mathcal{C}} u_c(\mathcal{A}) \geq \gamma, \quad \text{where } \mathcal{A} \in \mathcal{A}^*,$$

where the left term is the utility of the worst-off community and γ is the highest value for which the constraint is feasible.

Diversity Constraints (DC). Inspired by the game theoretic notion of core, DC requires that every community obtains a

³Our framework is also applicable to other forms of diffusion such as Linear Threshold Model (Kempe, Kleinberg, and Tardos 2003)

utility higher than when it receives resources proportional to its size and allocates them internally (Tsang et al. 2019). This is illustrated by the following constraint where U_c denotes the maximum utility that community c can achieve with a budget equal to $\lfloor KN_c/N \rfloor$.

$$u_c(\mathcal{A}) \geq U_c, \quad \forall c \in \mathcal{C} \text{ where } \mathcal{A} \in \mathcal{A}^*. \quad (2)$$

DC sets utility lower bounds for communities based on their relative sizes and how well they can spread influence internally. As a result, it does not explicitly account for reducing inequalities and may lead to high influence gap. We show this both theoretically and empirically in Sections and .

Demographic Parity (DP). Formalizing the legal doctrine of disparate impact (Zafar et al. 2017), DP requires the utility of all communities to be roughly the same. For any $\delta \in [0, 1)$, DP implies the constraints (Ali et al. 2019; Stoica, Han, and Chaintreau 2020; Aghaei, Azizi, and Vayanos 2019)

$$|u_c(\mathcal{A}) - u_{c'}(\mathcal{A})| \leq \delta, \quad \forall c, c' \in \mathcal{C} \text{ where } \mathcal{A} \in \mathcal{A}^*.$$

The degree of tolerated inequality is captured by δ and higher δ values are associated with higher tolerance. We use exact and approximate DP to distinguish between $\delta = 0$ and $\delta > 0$.

Fair Influence Maximization

Cardinal Welfare Theory Background

Following the cardinal welfare theory (Bergson 1938), our aim is to design welfare functions to measure the goodness of the choice of influencers. Cardinal welfare theory proposes a set of principles and welfare functions that are expected to satisfy these principles. Given two utility vectors, these principles determine if they are indifferent or one of them is preferred. For ease of exposition, let W denote this welfare function defined over the utilities of all individuals in the population (we will formalize W shortly). Then the existing principles of social welfare theory can be summarized as follows. Throughout this section, without loss of generality, we assume all utility vectors belong to $[0, 1]^N$.

(1) Monotonicity. If $\mathbf{u} \prec \mathbf{u}'$, then $W(\mathbf{u}) < W(\mathbf{u}')$.⁴ In other words, if \mathbf{u}' Pareto dominates \mathbf{u} , then W should strictly prefer \mathbf{u}' to \mathbf{u} . This principle also appears as *levelling down* objection in political philosophy (Parfit 1997).

(2) Symmetry. $W(\mathbf{u}) = W(P(\mathbf{u}))$, where $P(\mathbf{u})$ is any element-wise permutation of \mathbf{u} . According to this principle, W does not depend on the naming or labels of the individuals, but only on their utility levels.

(3) Independence of Unconcerned Individuals. Let $(\mathbf{u}|^c b)$ be a utility vector that is identical to \mathbf{u} , except for the utility of individual c which is replaced by a new value b . The property requires that for all c, b, b', \mathbf{u} and \mathbf{u}' , $W(\mathbf{u}|^c b) < W(\mathbf{u}'|^c b) \Leftrightarrow W(\mathbf{u}|^c b') < W(\mathbf{u}'|^c b')$. Informally, this principle states that W should be independent of individuals whose utilities remain the same.

(4) Affine Invariance. For any $a > 0$ and b , $W(\mathbf{u}) < W(\mathbf{u}') \Leftrightarrow W(a\mathbf{u} + b) < W(a\mathbf{u}' + b)$ i.e., the relative ordering is invariant to a choice of numeraire. **(5) Transfer Principle (Dalton 1920; Pigou 1912).** Consider individuals i and j in utility vector \mathbf{u} such that $u_i < u_j$. Let \mathbf{u}' be

⁴ \prec means $\mathbf{u}_c \leq \mathbf{u}'_c$ for all $c \in \mathcal{C}$ and $\mathbf{u}_c < \mathbf{u}'_c$ for some $c \in \mathcal{C}$.




Figure 1: The effect of network structure and in particular between-community edges on coupling of the utilities of communities. The figure shows two sample networks consisting of three communities, differentiated by shape: (a) is the same as (b) except that between-community edges are removed. Black fillings show the choice of influencers. We further assume p is small enough such that influence spread dissipates after one step. Transferring an influencer from circles to squares (top to bottom panel) affects the utility of diamonds in (b) but not in (a).

another utility vector that is identical to \mathbf{u} in all elements except i and j where $u'_i = u_i + \delta$ and $u'_j = u_j - \delta$ for some $\delta \in (0, (u_j - u_i)/2)$. Then, $W(\mathbf{u}) < W(\mathbf{u}')$. Informally, transferring utility from a high-utility to a low-utility individual should increase social welfare.

It is well-known that any welfare function W that satisfies the first four principles is additive and in the form of $W_\alpha(\mathbf{u}) = \sum_{i=1}^N u_i^\alpha / \alpha$ for $\alpha \neq 0$ and $W_\alpha(\mathbf{u}) = \sum_{i=1}^N \log(u_i)$ for $\alpha = 0$. Further, for $\alpha < 1$ the last principle is also satisfied. In this case α can be interpreted as an inequality aversion parameter, where smaller α values exhibit more aversion towards inequalities. We empirically investigate the effect of α in Section .

Group Fairness and New Principles

Applying the cardinal social welfare framework to influence maximization problems comes with new challenges. We next highlight these challenges and demonstrate our approach.

First, the original framework of cardinal welfare theory defines the welfare function over individuals. This is equivalent to seeking equality in the probability that each individual will be influenced, similar to the work of Fish et al. (2019). It is notoriously hard to achieve individual fairness in practice, e.g., Dwork et al. (2012) explore this in the machine learning context. The problem is further exacerbated in influence maximization because it is not always possible to spread the influence to isolated or poorly connected individuals effectively. Therefore, we focus on group fairness whereby the utility of each individual is defined as the average utility of the members of that community. Let u_c denote the average utility of community c . With this group-wise view, a welfare function can be written in terms of the average utilities over communities e.g., $W_\alpha(\mathbf{u}) = \sum_{i=1}^N u_i^\alpha / \alpha = \sum_{c \in \mathcal{C}} N_c u_c^\alpha / \alpha$.

Moreover, while principles 1-4 can be easily extended to our influence maximization problem, this is not the case for the transfer principle. More precisely, in the influence maxi-

mization problem it might not be feasible to directly transfer utilities from one community to another without affecting the utilities of other communities. We highlight this effect with an example, see Figure 1. In this figure, each community is represented by a distinct shape. The two networks (a) and (b) are identical except that between community edges are removed in network (a) (i.e., disconnected communities). The solid black vertices determine the choice of influencers. In network (b), if we transfer an influencer vertex from *circles* to *squares* according to Figure 1 (top to bottom panel), it will indirectly affect *diamonds* as well. This effect is absent in network (a) as there are no between-community edges to allow the spread of influence across communities. In network (a), the transfer principle prefers the resulting utility vector after the transfer. However, this principle cannot be applied to network (b) as the utilities of more than one community is modified after the transfer. Additionally, even when direct transfer is possible, it can be the case that there is no symmetry in the amount of utility gained by low-utility community and the amount of utility lost by high-utility community after the transfer. To address both of these shortcomings we introduce the *influence transfer principle* as a generalization of the transfer principle for influence maximization problems.

Similar to the original transfer principle, we consider solutions in which influencer vertices are transferred from one community to another community. Without loss of generality, we focus on the case where only one influencer vertex is transferred between the two communities. We refer to such solutions as neighboring solutions. Clearly, transfer of more than one influencer vertex can be seen as a sequence of transfers between neighboring solutions.

(5) Influence Transfer Principle. Let \mathcal{A} and $\mathcal{A}' \in \mathcal{A}^*$ be two neighboring solutions with corresponding utility vectors $\mathbf{u} = \mathbf{u}(\mathcal{A})$ and $\mathbf{u}' = \mathbf{u}(\mathcal{A}')$. Suppose the elements of \mathbf{u} and \mathbf{u}' are sorted in ascending order. We also assume after the transfer, the ordering of the utilities stays the same across \mathbf{u} and \mathbf{u}' .

If $\sum_{\kappa \in \mathcal{C}: \kappa \leq c} N_{\kappa}(u'_{\kappa} - u_{\kappa}) \geq 0 \forall c \in \mathcal{C}$ and $u'_c > u_c$ for some $c \in \mathcal{C}$, then $W(\mathbf{u}) < W(\mathbf{u}')$.

Informally, influence transfer principle states that in a desirable transfer of utilities, the magnitude of the improvement in lower-utility communities should be at least as high as the magnitude of decay in higher-utility communities while enforcing that at least one low-utility community receives a higher utility after the transfer. The original transfer is a special case of influence transfer principle when communities are disconnected and utilities transferred remain the same.

Next, we study whether any of the welfare functions that satisfy the first 4 principles satisfy the influence transfer principle. In Proposition 1, we show any additive and strictly concave function satisfies influence transfer principle. Since functions that satisfy the first 4 principles are strictly concave for $\alpha < 1$, the influence transfer principle is automatically satisfied in this regime. We defer all proofs to the full version.

Proposition 1. *Any strictly concave and additive function satisfies the influence transfer principle.*

To measure inequality, notion of utility gap (or analogous notions such as ratio of utilities) is commonly used (Fish

et al. 2019; Stoica, Han, and Chaintreau 2020). Utility gap measures the difference between the utilities of a pair of communities. In this work, we focus on the maximum utility gap, i.e., the gap between communities with the highest and lowest utilities (utility gap henceforth). For a utility vector \mathbf{u} , we define $\Delta(\mathbf{u}) = \max_{c \in \mathcal{C}} u_c - \min_{c \in \mathcal{C}} u_c$ to denote the utility gap. Fair interventions are usually motivated by the large utility gap before the intervention (Marmot et al. 2008). Fish et al. (2019) has shown that in social networks the utility gap can further increase after an algorithmic influence maximizing intervention. We extend this result to the entire class of welfare functions that we study in this work and we notice that the utility gap can increase even if we optimize for these welfare functions. This is a surprising result since, unlike the influence maximization objective, these welfare functions are designed to incorporate fairness, yet we may observe an increase in the utility gap. We now introduce another principle which aims to address this issue. Again we focus on neighboring solutions.

(6) Utility Gap Reduction. Let \mathcal{A} and $\mathcal{A}' \in \mathcal{A}^*$ be two neighboring solutions with corresponding utility vectors $\mathbf{u} = \mathbf{u}(\mathcal{A})$ and $\mathbf{u}' = \mathbf{u}(\mathcal{A}')$. If $\sum_{c \in \mathcal{C}} N_c u_c \leq \sum_{c \in \mathcal{C}} N_c u'_c$ and $\Delta(\mathbf{u}) > \Delta(\mathbf{u}')$ then $W(\mathbf{u}) < W(\mathbf{u}')$.

The utility gap reduction simply states that the welfare function should prefer the utility vector whose total utility is at least as high as the other vector and also has smaller utility gap. We now show that, in general, it is not possible to design a welfare function that obeys the utility gap reduction principle along with the other principles.

Proposition 2. *Let W be a welfare function that obeys principles 1-5. Then there exists an instance of influence maximization where W does not satisfy the utility gap reduction.*

Next, we show on a special class of networks, i.e., networks with disconnected communities, the utility gap reduction principle is satisfied in all influence maximization problems.

Proposition 3. *Let W be a welfare function that obeys principles 1-5. If the communities are disconnected, then W also satisfies the utility gap reduction principle.*

Propositions 2 and 3 and their proofs establish new challenges in fair influence maximization. These challenges arise due to the coupling of the utilities as a result of the network structure and more precisely the between-community edges. The results in Propositions 2 and 3 leave open the following question: “In what classes of networks, there exists a welfare function that satisfies all the 6 principles over all instances of influence maximization problems?” As an attempt to answer this question, we empirically show that over various real and synthetic networks including stochastic block models, there exist welfare functions that obey all of our principles. We conclude this section by the following three remarks.

Remark 1 (Application to Other Settings). *Our welfare-based framework can be theoretically applied to different graph-based problems (e.g., facility location) but algorithmic solution is domain-dependent. The choice of influence maximization is motivated by evidence about discrimination studied in previous work (Rahmatalabi et al. 2019; Stoica, Han, and Chaintreau 2020).*

Notion/Principle	Monotonicity	Symmetry	Ind. of Unconcerned	Affine	Influence Transfer	Gap Reduction
Exact DP	✗ (Prop. 5)	✓	✗ (Prop. 8)	✓	✗ (Prop. 11)	✓ (Prop. 15)
Approximate DP	✗ (Prop. 6)	✓	✗ (Prop. 8)	✗	✗ (Prop. 11)	✗ (Prop. 15)
DC	✓ (Cor. 1)	✗	✗ (Prop. 9)	✗	✗ (Prop. 12)	✗ (Prop. 16)
MMF	✓ (Cor. 1)	✓	✗ (Prop. 10)	✓	✗ (Prop. 13)	✗ (Prop. 17)
Utilitarian	✓ (Cor. 1)	✓	✓	✓	✗ (Prop. 14)	✗ (Prop. 18)
Welfare Theory	✓	✓	✓	✓	✓	✗ (Prop. 2)

Table 1: Summary of the properties of different fairness notions through the lens of welfare principles for influence maximization.

Remark 2 (Relationship between Principles and Fairness). *Monotonicity ensures there is no wastage of utilities. Symmetry enforces the decision-maker to not discriminate based on communities' names. According to the Independence of Unconcerned Individuals, between two solutions (choices of influencers) only those individuals/communities whose utilities change should impact the decision-maker's preference. Affine Invariance is a natural requirement that the preferences over different solutions should not change based on the choice of numeraire. Finally, the Transfer Principle promotes solutions that are more equitable.*

Remark 3 (Selecting the Inequality Aversion Parameter in Practice). *In our approach, α is a user-selected parameter that the user can vary to tune the trade-off between efficiency and fairness. Leaving the single parameter α in the hands of the user is a benefit of our approach since the user can inspect the solution as α is varied to select their preferred solution. Since a single parameter must be tuned, this can be done without the need for a tailored algorithm. In particular, we recommend that α be either selected by choosing among a moderate number of values and picking the one with the most desirable behavior for the user or by using the bisection method. Typically, choosing α will reduce to letting the user select how much utility gap they are willing to tolerate: they will select the largest possible value of α for which the utility gap is acceptable.*

Group Fairness and Welfare Maximization

The welfare principles reflect the preferences of a fair decision-maker between a pair of solutions. Thus a welfare function that satisfies all the principles would always rank the preferred (in terms of fairness and efficiency) solution higher. As a result, we can maximize the welfare function to get the most preferred solution.

We show that the two classes of welfare functions $W_\alpha(\mathbf{u}) = \sum_{i=1}^N u_i^\alpha / \alpha$ for $\alpha < 1, \alpha \neq 0$ and $W_\alpha(\mathbf{u}) = \sum_{i=1}^N \log(u_i)$ for $\alpha = 0$ satisfy 5 of our principles. Hence as a natural notion of fairness we can define a fair solution to be a choice of influencers with the highest welfare as defined in the following optimization problem.

$$\underset{\mathcal{A} \in \mathcal{A}^*}{\text{maximize}} \quad W_\alpha(\mathbf{u}(\mathcal{A})). \quad (3)$$

Lemma 1. *In the influence maximization problem, any welfare function that satisfies principles 1-5 is monotone and submodular.*

It is well-known that to maximize any monotone submodular function, there exists a greedy algorithm with a $(1 - 1/e)$ approximation factor (Nemhauser and Wolsey 1981) which we can also use to solve the welfare maximization problem.

Each choice of the inequality aversion parameter α results in a different welfare function and hence a fairness notion. A decision-maker can directly use these welfare functions as objective of an optimization problem and study the trade-off between fairness and total utility by varying α , see Section 5.

Connection to Existing Notions of Fairness

Our framework allows for a spectrum of fairness notions as a function of α . It encompasses as a special case *leximin fairness*⁵, a sub-class of MMF, for $\alpha \rightarrow -\infty$. Proportional fairness (Bonald and Massoulié 2001), a notion for resource allocation problems, is also closely connected to the welfare function for $\alpha = 0$. See the full version for more details.

It is natural to ask which of the fairness principles are satisfied by the existing notions of fairness for influence maximization. As we discussed in Section , the existing notions of fairness are imposed by adding constraints to the influence maximization problem. However, our welfare framework directly incorporates fairness into the objective. In order to facilitate the comparison, instead of the constrained influence maximization problems we consider an equivalent reformulation in which we bring the constraints into the objective via the characteristic function of the feasible set. We then have a single objective function which we can treat as the welfare function corresponding to the fairness constrained problem. More formally, given an influence maximization problem and fairness constraints written as a feasible set \mathcal{F}

$$\max_{\mathcal{A} \in \mathcal{A}^*} \sum_{c \in \mathcal{C}} N_c u_c(\mathcal{A}) \quad \text{s.t.} \quad \mathbf{u}(\mathcal{A}) \in \mathcal{F}.$$

We consider the following equivalent optimization problem

$$\max_{\mathcal{A} \in \mathcal{A}^*} \sum_{c \in \mathcal{C}} N_c u_c(\mathcal{A}) + \mathcal{I}_{\mathcal{F}}(\mathbf{u}(\mathcal{A})) := \max_{\mathcal{A} \in \mathcal{A}^*} W_{\mathcal{F}}(\mathbf{u}(\mathcal{A})),$$

in which $\mathcal{I}_{\mathcal{F}}(\mathbf{u})$ is equal to 0 if $\mathbf{u} \in \mathcal{F}$ and $-\infty$ otherwise. Using this new formulation, we can now examine each of the existing notions of fairness through the lens of the welfare principles. Given the new interpretation, to show that a fairness notion does not satisfy a specific principle, it suffices

⁵Leximin is subclass of MMF. According to its definition, among two utility vectors, leximin prefers the one where the worst utility is higher. If the worst utilities are equal, leximin repeats this process by comparing the second worst utilities and so on.

to show there exist solutions $\mathcal{A}, \mathcal{A}' \in \mathcal{A}^*$ and corresponding utility vectors $\mathbf{u} = \mathbf{u}(\mathcal{A})$ and $\mathbf{u}' = \mathbf{u}(\mathcal{A}')$ such that the principle prefers \mathbf{u} over \mathbf{u}' but $W_{\mathcal{F}}(\mathbf{u}) < W_{\mathcal{F}}(\mathbf{u}')$. The results are summarized in Table 1 where in addition to comparing with the previous notions introduced in Section , we compare with the utilitarian notion i.e., Problem (1). We provide formal proofs for each entry of Table 1 in the full version.

We observe that none of the previously defined notions of fairness for influence maximization satisfies all of our principles and each existing notion violates at least 3 out of 6 principles. We point out that exact DP is the only notion that satisfies the utility gap reduction. However, this comes at a cost as enforcing exact DP may result in significant reduction in total utility in the fair solution compared to the optimal unconstrained solution (Corbett-Davies et al. 2017).

Computational Results

We evaluate our approach in terms of both the total utility or spread of influence (to account for efficiency) and utility gap (to account for fairness). We show by changing the inequality aversion parameter, we can effectively trade-off efficiency with fairness. As baselines, we compare with DC and MMF. To the best of our knowledge, there is no prior work that handles DP constraints over the utilities. We follow the approach of Tsang et al. (2019) for both problems and view these problems as a multi-objective submodular optimization with utility of each each community being a separate objective. They propose an algorithm and implementation with asymptotic $(1 - 1/e)$ approximation guarantee which we also utilize here. We use *Price of Fairness* (PoF), defined as the percentage loss in the total influence spread as a measure of efficiency. Precisely, $\text{PoF} := 1 - \text{OPT}^{\text{fair}}/\text{OPT}^{\text{IM}}$ in which OPT^{fair} and OPT^{IM} are the the total influence spread, with and without fairness. Hence $\text{PoF} \in [0, 1]$ and smaller values are more desirable. The normalization in PoF allows for a meaningful comparison between networks with different sizes and budgets as well as between different notions of fairness. In the PoF calculations, we utilize the generic greedy algorithm (Kempe, Kleinberg, and Tardos 2003) to compute OPT^{IM} . To account for fairness, we compare the solutions in terms of the utility gap. Analogous measures are widely used in fairness literature (Hardt, Price, and Srebro 2016) and more recently in graph-based problems (Fish et al. 2019; Stoica, Han, and Chaintreau 2020). We also note that our framework ranks solutions based on their welfare and does not directly optimize utility gap, as such our evaluation metric of fairness does not favor any particular approach.

We perform experiments on both synthetic and real networks. We study two applications: community preparedness against landslide incidents and suicide prevention among homeless youth. We discuss the latter in the full version. In the synthetic experiments, we use the *stochastic block model* (SBM) networks, a widely used model for networks with community structure (Fienberg and Wasserman 1981). In SBM networks, vertices are partitioned into disjoint communities. Within each community c , an edge between two vertices is present independently with probability q_c . Between any two vertices in communities c and c' , an edge exists indepen-

dently with probability $q_{cc'}$ and typically $q_c > q_{cc'}$ to capture homophily (McPherson, Smith-Lovin, and Cook 2001). SBM captures the community structure of social networks (Girvan and Newman 2002). We report the average results over 20 random instances and set $p = 0.25$ in all experiments.

Landslide Risk Management in Sitka, Alaska. Sitka, Alaska is subject to frequent landslide incidents. In order to improve communities' preparedness, an effective approach is to instruct people on how to protect themselves before and during landslide incidents. Sitka has a population of more than 8000 and instructing everyone is not feasible. Our goal is to select a limited set of individuals as peer-leaders to spread information to the rest of the city. The Sitka population is diverse including different age groups, political views, seasonal and stable residents where each person can belong to multiple groups. These groups differ in their degree of connectedness. This makes it harder for some groups to receive the intended information and also impacts the cost of imposing fairness.

Since collecting the social network data for the entire city is cumbersome, we assume a SBM network and use in-person semi-structured interview data from 2018-2020 with members of Sitka to estimate the SBM parameters. Using the interview responses in conjunction with the voter lists, we identified 5940 individuals belonging to 16 distinct communities based on the intersection of age groups, political views, arrival time to Sitka (to distinguish between stable and transient individuals). The size of the communities range from 112 (stable, democrat and 65+ years of age) to 693 (republican, transient fishing community, age 30-65). See the full version for details on the estimation of network parameters.

Figure 2 summarizes results across different budget values K ranging from 2% to 10% of the network size N for our framework (different α values) as well as the baselines. In the left panel, we observe that as α decreases, our welfare-based framework further reduces the utility gap, achieving lower gap than DC and competitive gap as MMF. As we noted in Section , our framework recovers leximin (which has stronger guarantees than MMF) as $\alpha \rightarrow -\infty$, though we show experimentally that this is achieved with moderate values of α . Overall, utility gap shows an increasing trend with budget, however the sensitivity to budget decreases when more strict fairness requirements are in place, e.g. in MMF and $\alpha = -9.0$. From the right panel, PoF varies significantly across different approaches and budget values surpassing 40% for MMF. This is due to the stringent requirement of MMF to raise the utility of the worst-off as much as possible. Same holds true for lower values of α as they exhibit higher aversion to inequality. The results also indicate that PoF decreases as K grows which captures the intuition that fairness becomes less costly when resources are in greater supply. Resource scarcity is true in many practical applications, including the landslide risk management domain which makes it crucial for decision-makers to be able to study different fairness-efficiency trade-offs to come up with the most effective plan. Figure 3 depicts such trade-off curves where each line corresponds to a different budget value across the range of α . Previous work only allows a decision-maker to choose among a very limited set of fairness notions regardless of the application requirements. Here, we show that our framework




Figure 2: Left and right panels: utility gap and PoF for different K and α values for our framework and baselines.

allows one to choose α to meaningfully study the PoF-utility gap trade-offs. For example, given a fixed budget and a tolerance on utility gap, one can choose an α with the lowest PoF. We now investigate the effect of relative connectedness. We study the effect of relative community size in the full version.

Relative Connectedness. We sample SBM networks consisting of 3 communities each of size 100 where communities differ in their degree of connectedness. We set $q_1 = 0.06, q_2 = 0.03, q_3 = 0.0$ to obtain three communities with high, mid and low relative connectedness. We choose these values to reflect asymmetry in the structure of different communities which mirrors real world scenarios since not every community is equally connected. We set between-community edge probabilities $q_{cc'}$ to 0.005 for all c and c' and $K = 0.1N$. We gradually increase q_3 from 0.0 to 0.06. Results are summarized in Figure 4, where each group of bars correspond to a different approach. We observe as q_3 increases utility gap and PoF decrease until they reach a minimum around at around $q_3 = 0.03$. From this point, the trend reverses. This U-shaped behavior is due to structural changes in network. More precisely, for $q_3 < 0.03$ we are in the high-mid-low connectedness regime for the three groups, where the third community receives the minimum utility. As a result, as q_3 increases it becomes more favorable to choose more influencer vertices in this community which in turn reduces the utility gap. For $q_3 > 0.03$, the second community will be become the new worst-off community due its lowest connectedness. Hence, further increase in q_3 causes more separation in connectedness and we see previous behavior in reverse. Thus, by further increasing q_3 , communities 1 and 3 receive more and more influencer vertices. This behavior

translates to PoF as the relative connectedness of communities impacts how *hard* it is to achieve a desired level of utility gap. Finally, we see that the U-shaped behavior is skewed, i.e., we observe higher gap and PoF in lower range of q_3 which is due to higher gap in connectedness of communities. We can also compare the effect of relative connectedness and community size (see the full version.) We observe that connectedness has a more significant impact on PoF (up to 25%) compared to community size (less than 4%). In other words, when communities are structurally different it is more costly to impose fairness. This is an insightful result given that in different applications we may encounter different populations with structurally different networks. Utility gap on the other hand is affected by both size and connectedness. Finally while our theory indicates that in the network setting, no welfare function can satisfy all principles including utility gap reduction over all instances of the influence maximization, we observe that our class of welfare functions satisfies all of the desiderata on the class of networks that we empirically study. Our theoretical results showed this for a special case of networks with disconnected communities. In particular, we see higher PoF is accompanied by lower utility gap which complies with utility gap reduction principle.




Figure 3: PoF vs. utility gap trade-off curves. Each line corresponds to a different budget K across different α values.




Figure 4: Utility gap and PoF for various levels q_3 . All results are compared across different values of α and the baselines.

Broader Impact

As the empirical evidence highlighting ethical side effects of algorithmic decision-making is growing (Angwin et al. 2016; Miller 2015), the nascent field of algorithmic fairness has also witnessed a significant growth. It is well-established by this point that there is no universally agreed-upon notion of fairness, as fairness concerns vary from one domain to another (Narayanan 2018; Berk et al. 2017). The need for different fairness notions can also be explained by theoretical studies that show that different fairness definitions are often in conflict with each other (Kleinberg, Mullainathan, and Raghavan 2017; Chouldechova 2017; Friedler, Scheidegger, and Venkatasubramanian 2016). To this end, most of the literature on algorithmic fairness proposes different fairness notions motivated by different ethical concerns. A major drawback of this approach is the difficulty of comparing these methods against each other in a systematic manner to choose an appropriate notion for the domain of interest. Instead of following this trend, we propose a unifying framework controlled by a single parameter that can be used by a decision-maker to systematically compare different fairness measures which typically result in different (and possibly also problem-dependent) trade-offs. Our framework also accounts for the social network structure while designing fairness notions – a consideration that is mainly overlooked in the past. Given these two contributions, it is perceivable that our approach can be used in many of the public health interventions such as suicide, HIV or Tuberculous prevention that rely on social networks. This way, the decisions-makers can compare a menu of fairness-utility trade-offs proposed by our approach and decide which one of these trade-offs are more desirable without a need to understand the underlying mathematical details that are used in deriving these trade-offs.

There are crucial considerations when deploying our system in practice. First, cardinal welfare is one particular way of formalizing fairness considerations. This by no means implies that other approaches for fairness e.g. equality enforcing interventions should be completely ignored. Second, we have assumed that the decision-maker has the full knowledge of the network as well as possibly protected attributes of the individuals which can be used to define communities. Third, while our experimental evaluation is based on utilizing a greedy algorithm, it is conceivable that this greedy approximation can create complications by imposing undesirable biases that we have not accounted for. Intuitively (and as we have seen in our experiments) the extreme of inequality aversion ($\alpha \rightarrow -\infty$) can be used as a proxy for pure equality. However, the last two concerns require more care and we leave the study of such questions as future work.

Acknowledgments

We would like to thank David Gray Grant, Matthew Joseph, Andrew Perrault and Bryan Wilder for helpful discussions about this work. This work is supported in part by the Smart & Connected Communities program of the National Science Foundation under NSF award No. 1831770, the US Army Research Office under grant number W911NF1710445 and Center for Research on Computation and Society (CRCS)

at the Harvard John A. Paulson School of Engineering and Applied Sciences.

References

- Abasolo, I.; and Tsuchiya, A. 2004. Exploring social welfare functions and violation of monotonicity: an example from inequalities in health. *Journal of Health Economics* 23(2): 313–329.
- Aghaei, S.; Azizi, M. J.; and Vayanos, P. 2019. Learning optimal and fair decision trees for non-discriminative decision-making. In *33rd AAAI Conference on Artificial Intelligence*, volume 33, 1418–1426.
- 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 abs/1905.06618*.
- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2016. Machine Bias. *ProPublica*.
- Bergson, A. 1938. A Reformulation of Certain Aspects of Welfare Economics. *The Quarterly Journal of Economics* 52(2): 310–334.
- Berk, R.; Heidari, H.; Jabbari, S.; Joseph, M.; Kearns, M.; Morgenstern, J.; Neel, S.; and Roth, A. 2017. A Convex Framework for Fair Regression. In *4th Workshop on Fairness, Accountability, and Transparency in Machine Learning*.
- Bonald, T.; and Massoulié, L. 2001. Impact of fairness on Internet performance. In *Joint International Conference on Measurements and Modeling of Computer Systems*, 82–91.
- Chouldechova, A. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. In *3rd Workshop on Fairness, Accountability, and Transparency in Machine Learning*.
- Corbett-Davies, S.; Pierson, E.; Feller, A.; Goel, S.; and Huq, A. 2017. Algorithmic decision making and the cost of fairness. In *23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 797–806.
- Dalton, H. 1920. The Measurement of the Inequality of Incomes. *The Economic Journal* 30(119): 348–361.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *3rd Innovations in Theoretical Computer Science*, 214–226.
- Fienberg, S.; and Wasserman, S. 1981. Categorical Data Analysis of Single Sociometric Relations. *Sociological Methodology* 12: 156–192.
- Fish, B.; Bashardoust, A.; Boyd, D.; Friedler, S.; Scheidegger, C.; and Venkatasubramanian, S. 2019. Gaps in Information Access in Social Networks? In *28th World Wide Web Conference*, 480–490.
- Friedler, S.; Scheidegger, C.; and Venkatasubramanian, S. 2016. On the (im)possibility of fairness. *arXiv abs/1609.07236*.
- Girvan, M.; and Newman, M. E. 2002. Community structure in social and biological networks. *Proceedings of the national academy of sciences* 99(12): 7821–7826.

- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. In *Advances in neural information processing systems* 29, 3315–3323.
- Heidari, H.; Ferrari, C.; Gummadi, K.; and Krause, A. 2018. Fairness behind a veil of ignorance: A welfare analysis for automated decision making. In *Advances in Neural Information Processing Systems* 31, 1265–1276.
- Isaac, M.; Elias, B.; Katz, L.; Belik, S.-L.; Deane, F.; Enns, M.; Sareen, J.; and Team, S. C. S. P. 2009. Gatekeeper training as a preventative intervention for suicide: a systematic review. *The Canadian Journal of Psychiatry* 54(4): 260–268.
- Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 137–146.
- Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2017. Inherent Trade-Offs in the Fair Determination of Risk Scores. In *8th Conference on Innovations in Theoretical Computer Science*, 43:1–43:23.
- Marmot, M.; Friel, S.; Bell, R.; Houweling, T. A.; Taylor, S.; on Social Determinants of Health, C.; et al. 2008. Closing the gap in a generation: health equity through action on the social determinants of health. *The lancet* 372(9650): 1661–1669.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* 27(1): 415–444.
- Miller, C. 2015. Can an Algorithm Hire Better than a Human? *The New York Times* URL <http://www.nytimes.com/2015/06/26/upshot/can-an-algorithm-hire-better-than-a-human.html/>. Ret. 4/28/2016.
- Narayanan, A. 2018. Translation tutorial: 21 fairness definitions and their politics. In *Proc. Conf. Fairness Accountability Transp., New York, USA*, volume 2, 6–2.
- Nemhauser, G. L.; and Wolsey, L. A. 1981. Maximizing submodular set functions: formulations and analysis of algorithms. In *North-Holland Mathematics Studies*, volume 59, 279–301. Elsevier.
- Parfit, D. 1997. Equality and priority. *Ratio* 10(3): 202–221.
- Pigou, A. 1912. *Wealth and welfare*. Macmillan and Company.
- Rahmatalabi, 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* 32, 15750–15761.
- Rawls, J. 2009. *A theory of justice*. Harvard University Press.
- Sen, A. 1997. *On economic inequality*. Oxford university press.
- Stoica, A.-A.; Han, J. X.; and Chaintreau, A. 2020. Seeding Network Influence in Biased Networks and the Benefits of Diversity. In *Proceedings of The Web Conference*, 2089–2098.
- Tsang, A.; Wilder, B.; Rice, E.; Tambe, M.; and Zick, Y. 2019. Group-Fairness in Influence Maximization. In *28th International Joint Conference on Artificial Intelligence*, 5997–6005.
- 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.
- Wilder, B.; Onasch-Vera, L.; Diguseppi, G.; Petering, R.; Hill, C.; Yadav, A.; Rice, E.; and Tambe, M. 2020. Clinical trial of an AI-augmented intervention for HIV prevention in youth experiencing homelessness.
- Yonemoto, N.; Kawashima, Y.; Endo, K.; and Yamada, M. 2019. Gatekeeper training for suicidal behaviors: A systematic review. *Journal of affective disorders* 246: 506–514.
- Zafar, M. B.; Valera, I.; Gomez-Rodriguez, M.; and Gummadi, K. 2017. Fairness Constraints: Mechanisms for Fair Classification. In *20th International Conference on Artificial Intelligence and Statistics*, 962–970.