

Intervening to Increase Community Trust for Fair Network Outcomes

Naina Balepur
nainab2@illinois.edu
University of Illinois Urbana-Champaign
Urbana, IL, USA

Hari Sundaram
hs1@illinois.edu
University of Illinois Urbana-Champaign
Urbana, IL, USA

ABSTRACT

Refugees or immigrants who arrive in new countries often feel isolated. In this work, we examine how a resource-bounded public entity can make recommendations to increase integration of these new arrivals into a community. The community is made up of agents who engage in a strategic network formation process; agents join periodically — new arrivals are the refugees. The public entity meanwhile makes a limited number of edge-formation recommendations (according to its resource constraint) per iteration in order to increase integration of refugees. This work investigates the relationship between community trust and network fairness. First, we show that increasing the public entity’s resource allocation will *not* compensate for low trust in the community. Then, we introduce two trust-increasing interventions by the public entity: a targeted advertising campaign, and an announcement to increase transparency. We find that diverting a fraction (20%) of the public entity’s resources to a targeted advertising campaign can increase trust and fairness in the community, especially in low trust scenarios. We find that personalized, local announcements are more effective at increasing fairness than global announcements in low trust scenarios; they almost double our fairness metric in some cases. Importantly, the transparent announcement requires no extra resource expenditure on the part of the public entity. Our work underscores the importance of community trust — low trust cannot be compensated for with resources. This work provides theoretical support for these trust-increasing interventions, which we show can lead to increased integration of refugees in communities.

CCS CONCEPTS

- **Human-centered computing** → **Social recommendation**;
- **Computing methodologies** → **Agent / discrete models**;
- **Theory of computation** → **Social networks**.

KEYWORDS

trust, fairness, refugee resettlement, network formation

ACM Reference Format:

Naina Balepur and Hari Sundaram. 2024. Intervening to Increase Community Trust for Fair Network Outcomes. In *The 2024 ACM Conference on*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAccT '24, June 03–06, 2024, Rio de Janeiro, Brazil

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0450-5/24/06

<https://doi.org/10.1145/3630106.3659008>

Fairness, Accountability, and Transparency (FAccT '24), June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3630106.3659008>

1 INTRODUCTION

When seeking asylum in a new country, a refugee’s basic needs for food, housing, and employment are usually handled by governmental programs [35]. However, without explicit focus on social integration, they may begin to feel isolated. Attempting to cure this isolation without help is difficult; the refugee has entered a community where subgroups and cliques have already formed, perhaps many years ago. An intermediary like a nonprofit organization with expendable resources (such organizations are funded by the Office of Refugee Resettlement [36]) could help introduce the refugee to other community members to help them feel less isolated, if the refugee and community members are willing to *trust* this organization. In this work, we investigate the relationship between trust and community outcomes. We quantify how trust affects community integration for refugees, as well as the effectiveness of trust-increasing interventions.

Prior work on refugee relocation focuses most heavily on community matching [1, 16, 24]; none to our knowledge have considered the position of the refugee within the community. In fact, Nawyn [33] and Beiser [8, 9] underscore the need for social considerations by highlighting under-considered aspects of refugee resettlement: faith, ethnicity, culture, and mental health. Other research on the relationship between community and trust finds that participation in library programs increases trust in the library [45, 46]. These works give theoretical evidence that trust in a public entity *can* change with increased exposure. However, none investigate the more specific mechanisms for actually *creating* trust in these scenarios, nor do they study how increased trust can affect individual or community outcomes.

In this work, we develop a model where refugees enter a network and receive help from a fairness-minded, resource-bounded public entity. We begin with a set of original community members, and new agents join periodically; these are refugees or immigrants to the community. These agents strategically form edges with each other following an iterative network formation model. They are also initialized with some baseline trust in the public entity. New arrivals are at an innate disadvantage because they have less information regarding the network. The public entity attempts to decrease inequality caused by this disadvantage by making edge-formation recommendations to newly arrived agents. If an agent trusts the public entity, they will take the recommendation. Else, they will continue strategically forming edges.

The process described above is the baseline behavior of our model, which we use to simulate various interventions that might

increase trust, and therefore fairness. First, we ask whether increasing the public entity’s resources will compensate for low agent trust when trying to achieve fair community outcomes. When the average agent trust is at 0.29 or below (on a $[0, 1]$ scale), maximizing resources only allows the public entity to reach 7.86% of maximum community fairness, demonstrating the importance of trust regardless of resource allocation. Next, we explore a trust-increasing intervention: an advertising campaign targeted at agents who do not trust the public entity. We find that diverting a fraction (20%) of the public entity’s resources to a targeted advertising campaign can increase trust and fairness in the community, especially in lower community trust scenarios (under 0.40). The same is not true for larger fractions of resources (40% – 100%). Finally, we investigate whether a public entity that is transparent about its successes and failures would increase trust. We simulate two scopes: global and local. We find that personalized announcements regarding fairness in local neighborhoods are more effective in increasing trust and fairness than global announcements in low trust scenarios; they can almost double our fairness metric in some cases. Our contributions are as follows:

Refugee integration: We propose a framework that improves the network position of refugees in a network using edge formation recommendations from a public entity. Rather than focusing only on refugee community assignment [1, 16, 24], which ignores important social and cultural considerations [8, 9, 33], our work focuses on refugees’ network positions within these communities. Specifically, we model a resource-bounded public entity that makes recommendations in order to improve integration of new community members. We find that when trust in the public entity is high, it can have great impact on the community, even with relatively low resources.

Trust-increasing interventions: We develop two trust-increasing interventions and find that they are successful in increasing trust and fairness. In contrast, prior work has found correlations between trust and behavior [17, 44, 45], but we are not aware of any work that simulates trust-increasing interventions on communities. We find that 1) targeted advertising campaigns and 2) transparent public entities can successfully improve low trust, and by extension fairness. Transparency in particular is an important finding, as it doesn’t require any additional resources from the public entity to increase trust.

2 RELATED WORK

2.1 Refugee resettlement problem

Seminal theoretical work in this area focuses on mechanism design for refugee resettlement algorithms. Delacrétaz et al. [16] propose a refugee resettlement mechanism that builds off two-sided matching theory, considering both the desires of the individual and the community. Alternatively, Gözl and Procaccia [24] cast the problem as a submodular optimization task. Ahani et al. [1] develop a two-stage stochastic programming approach that optimizes employment for resettled employees. In fact, this algorithm is now used for refugee resettlement in the United States. Though these works present strong theoretical algorithmic results, they fail to consider political and social aspects. This can lead to problems which are highlighted

in other works. Nawyn [33] discusses aspects of refugee resettlement that may be ignored or under-considered: faith, ethnicity, and culture. Beiser [8] conducts a longitudinal study of Southeast Asian refugees in Canada, finding that mental health outcomes could be improved with specific interventions [9]. Other empirical works examine existing resettlement programs in various countries and propose adjustments [11, 21, 43, 50]. In this work we produce good theoretical results while also considering individual social benefits.

2.2 Fair networks

We develop our fair network definition by building off existing conceptions. Mehrotra et al. [31] survey fairness metrics usually used in machine learning, and find that ignoring relationships between individuals can introduce spurious fairness conclusions. They develop two frameworks for fairness on social networks; one captures how access to the network varies across groups, and the other captures inter-group biases. Liu et al. [30] develop a metric for group fairness on social networks that relies on homophily rather than sharing of protected demographic attributes. In the game theoretic context, Judd et al. [29] execute experiments to determine if human behavioral choices will lead to maximum social welfare for participants. Santos et al. [39] find that high structural power in networks leads to increasingly fair decision-making for the collective. Other works consider group and individual fairness when solving classic network problems. Saxena et al. [40] survey fairness in classic network problems broadly. Atwood et al. [6] study fair allocation for treatments such as vaccinations on a social network. Classic network topics like influence maximization and information flow can also be framed to consider fairness [3, 27, 28, 34, 38, 42]. For our problem, we build off prior work by developing a group fairness metric reliant only on final network structure. We utilize game theoretic results on fairness to design a mechanism where individual decisions likely increase network fairness.

2.3 Community trust

A key aspect of our model is trust from the community — community members must trust the public entity in order to take recommendations from it. Di Napoli et al. [19] conduct a study that uses community trust as an indicator for opportunities. They find that community trust is significantly associated with community engagement. Using a game theoretic perspective, Jachimowicz et al. [26] investigate whether poor individuals improve their myopic decision-making when community trust is increased. The researchers find that individuals in communities with higher levels of community trust make less myopic choices. This result is highly influential to our work; we design an algorithm where high community trust causes individuals to make decisions that account for long-term outcomes. Many contextual case studies have been conducted. In medicine, Webb Hooper et al. [49] find differing distrust levels in healthcare on the basis of race and disability status. Paton [37] discovers that situational factors and collective problem solving and empowerment play an important role in community trust for hazard preparedness. In work bridging knowledge sharing and community trust, Chen et al. [13] find that community trust impacts knowledge sharing intention and thus behavior in professional

teaching networks. Bridging community trust and refugee resettlement, Wallman Lundåsen and Wollebæk [47] study the effect of immigration-related diversity on different forms of trust. They find that asymmetry in norms and perceptions of unfairness negatively impact community trust, especially in diverse communities. These studies underscore the impact of community trust, but none experiment to find a causal link; our work bridges this gap.

2.4 Altruism

We can frame our public entity as an altruistic actor, but not in the usual sense. It has its own goals to achieve, but these goals will directly help individuals on the network in the long-term. Much work has been done on this topic. Andreoni [4], develops a model for altruism where agents receive a “warm glow” from giving. In another work [5], he develops the altruism coefficient, which splits an agent’s propensity to donate into egoistic (impure) and altruistic (pure) components. Fehr and Fischbacher [22] study the nature of human altruism. It has been shown that cooperation cannot survive even when there are some strong reciprocators in presence of many free-riders. DeSteno [18] hypothesizes that compassion is more likely shown to someone similar to you. Stevens and Hauser [41] propose that being altruistic is only non-beneficial in the short term; in the long run reciprocity may occur. Feiler et al. [23] propose that mixing egoistic and altruistic reasons reduces the likelihood of giving by increasing individuals’ awareness that a persuasion attempt is happening. Work has been done regarding altruism on networks specifically as well. Chen et al. [13] investigate the SCNet, a social network for teachers, to determine the role of altruism in knowledge sharing. Bourlès et al. [10] develop a model of altruism on networks. These agents care about the well-being of their neighbors and may provide financial support to their poorer friends. They show that a positive income shock to an individual weakly benefits all others. While we do not consider altruistic individual agents, we instead tackle a new and largely unexplored problem of an altruistic entity working externally to a network of selfish individuals.

3 PROBLEM STATEMENT

Consider an active community where individuals form and break connections with each other. The individuals in this community have been part of it since its inception, and are well informed regarding *who* is in the community and *how* they are connected. Now, consider new community members who join *after* a social network has formed. These newcomers are at a disadvantage when it comes to integration. To remedy this imbalance, a public entity exists to give edge formation (*i.e.*, friendship) recommendations to agents to become more integrated. Agents will take these recommendations if they trust the public entity. The problem addressed in this paper is inspired by the refugee resettlement problem, where refugees are resettled in communities on the basis of high employment probability. In this paper we examine the scenario in which a refugee has already been assigned to a community, and the public entity now wants to ensure proper social integration.

Formally, we start with a population of original community member V_O in a disconnected network $G = (V, E)$ where $V = V_O$. Each agent $v \in V$ has an attribute a_v and utility function f_v . The utility

function dictates which edges will give agent v positive utility, and is based on both agents’ attributes and local network structure. These agents $v \in V$ engage in an iterative network formation process \mathcal{N} , myopically choosing edges during each iteration i that give positive utility according to f_v . After some iterations, new agents V_R arrive to the network; we call this set of agents newcomers or refugees to match our intended application. Now, $G = (V, E)$ where $V = V_O \cup V_R$.

We now introduce the public entity P , which helps ensure fair community outcomes. The public entity P deploys recommendation algorithm \mathcal{A} , which takes social network G as input. Note that the algorithm \mathcal{A} does not have access to *all* network information; notably, the utility functions f_v of individuals $v \in V$ are not known by P . The public entity P also has limited capacity to give recommendations; it can give ρ -many per iteration i of network formation model \mathcal{N} .

The population of individuals V uses the public entity’s recommendation algorithm \mathcal{A} to make edge formation decisions in the following way. First, the public entity P will choose ρ -many agents in V via an agent selection process \mathcal{S} (recall that ρ is the public entity’s resource constraint). Then, P will provide each of these ρ -many agents a recommendation r_v dictated by algorithm \mathcal{A} . The recommendation r_v is given as a node u , which P recommends v form an edge with. The public entity makes recommendations which maximize its fairness goals \mathcal{F} . Each agent has a level of trust $\tau_v \in [0, 1]$ in the public entity, drawn from a distribution T . This trust represents the degree to which an agent v believes that the public entity’s recommendations will positively impact their community *in the long term*. As the fairness metric \mathcal{F} is a global measure, each agent v to some extent believes that global improvement will lead to individual improvement as well. Thus, with probability τ_v , agent v will blindly take recommendation r_v . With probability $1 - \tau_v$, v makes its utility-maximizing edge formation decision consistent with network formation model \mathcal{N} . Without a recommendation, agent v also continues with network formation process \mathcal{N} during iteration i .

In this paper we investigate the relationship between agents’ level of trust in the system and the public entity’s ability to achieve its policy goals \mathcal{F} . We ask three research questions:

- RQ1:** What is the trade-off between agent trust distribution T and the public entity’s resource constraint ρ ? Can increasing the public entity’s resources compensate for low agent trust as the public entity tries to improve \mathcal{F} every iteration?
- RQ2:** If an agent’s trust level τ_v is known by the public entity, can P target agents with low trust via advertising to increase trust and fairness \mathcal{F} ? How much money would P spend on this campaign?
- RQ3:** Consider a schema where the public entity P announces the effects of its interventions on fairness (negative or positive). If an agent’s trust level τ_v is affected by this announcement, can transparency from P lead to improved fairness \mathcal{F} ?

4 NETWORK FORMATION MODEL

We extend the network formation model developed by Christakis et al. [14] for the agents’ network formation process \mathcal{N} . We considered other canonical models: Erdős–Rényi [20], Barabási–Albert [2],

Watts-Strogatz [48], and the Stochastic Block Model [25]. However, these models do not involve strategic agents, and so do not fit our application.

In this section we give an overview of the network formation model, describing each key process: opportunities for establishing links, the link formation process, and agents' preferences. We start with a completely disconnected network of N -many agents at iteration $t = 0$, we call this set of agents V_O (original community members). We have m -many more agents enter one at a time starting at $t = 1$, these agents are set V_R (refugees). Each agent $u \in V = V_O \cup V_R$ is initialized with a categorical attribute α_u .

Opportunities for links: Every iteration, each agent $u \in V$ has the opportunity to form an edge. This opportunity comes uniformly at random from u 's second-degree neighbors. If none are available, a node is selected at random from $V \setminus u$.

Link formation: Once two nodes u and v are paired, they each calculate the utility of the edge from utility functions f_u and f_v respectively. If this edge does not yet exist, then each agent will agree to the edge formation if they both calculate a positive utility. If the edge already exists, then it will be dissolved if either agent experiences negative utility.

Agent preferences: The utility function f_u for each agent $u \in V$ is the same. Utility function f_u outputs the utility of edge (u, v) for agent u , which is a linear combination of the following: α_v , $\mathbb{1}(\alpha_u == \alpha_v)$ (i.e., whether u and v share the same attribute), v 's degree, the square of v 's degree, and whether u and v are separated by a distance of 2 or 3. These properties fully capture both the attributed and structural qualities of the connection. We use model fits from the original paper to choose reasonable coefficients for this utility function.

Algorithm 1 below gives the overview of our extension of Christakis's network formation process.

Algorithm 1: Network formation model, \mathcal{N}

Data: N the number of original agents, m the number of new agents, I the number of iterations
 // NETWORK INITIALIZATION
 1 $G \leftarrow (V, E)$ with $|V| = N$ and $E = \emptyset$;
 2 **for** $u \in V$ **do**
 3 $\alpha_u \leftarrow$ an attribute at random ;
 // NETWORK FORMATION
 4 **for** $i = 0 \dots I - 1$ **do**
 5 **if** $1 \leq i \leq m$ **then**
 6 $V \leftarrow V \cup v'$; // add new node
 7 **for** $u \in V$ **do**
 8 $edgeFormation(G, u)$
 9 **return** G

5 PUBLIC ENTITY MODEL

In this section we describe how the public entity P functions. We give an overview of the model, followed by subroutines and a formal algorithm. We also give details of our fairness metric in this section.

As agents follow the network formation model \mathcal{N} described in the previous section, the public entity selects a subset of agents via

Algorithm 2: edgeFormation

Data: $G = (V, E)$ the network, u the node forming an edge
 1 $v \leftarrow edgeOpportunity(G, u)$; // get rec
 2 **if** $(u, v) \notin E$ **then**
 3 **if** $f_u(u, v) > 0 \wedge f_v(u, v) > 0$ **then**
 4 $E \leftarrow E \cup (u, v)$; // add edge
 5 **else**
 6 **if** $f_u(u, v) < 0 \vee f_v(u, v) < 0$ **then**
 7 $E \leftarrow E \setminus (u, v)$; // remove edge

Algorithm 3: edgeOpportunity

Data: $G = (V, E)$ the network, u the node to pair up
 // $N_1(u)$ and $N_2(u)$ are the first and second degree neighbors of u respectively
 1 **if** $|N_2(u)| == 0$ **then**
 2 $v \leftarrow v \in V \setminus u$ uniformly at random ;
 3 **else**
 4 $v \leftarrow v \in N_2(u) \setminus (N_1(u) \cup u)$ uniformly at random ;
 5 **return** v

selection process \mathcal{S} . It then makes edge formation recommendations to these chosen agents using recommendation algorithm \mathcal{A} . The public entity P has a resource constraint ρ , so at most it can make ρ -many recommendations during one iteration of the network formation model.

Selection process \mathcal{S} : The selection process \mathcal{S} first produces an ordering of agents from highest to lowest priority, and then selects the first ρ -many of them. The ordering is done with the following prioritization: 1) new arrivals to the network in order of arrival 2) the original network members in a random order. This allows for new agents to get recommendations first, while allowing for original agents to get recommendations too if the resources allow for it. This prioritization aligns with our goal to help new agents V_R who have little information regarding the network. If there are more new arrivals than resources, they will not all receive recommendations.

Recommendation algorithm \mathcal{A} : After selecting agents, the public entity recommends an edge to each of them. The recommended edge myopically maximizes the fairness metric \mathcal{F} (i.e., community welfare). The public entity does not take into account all recommendations it makes during any given iteration; the maximization is always myopic. We make this choice because we assume that each interaction is long in duration; we expect agents to receive recommendations one at a time. Agents accept this recommendation blindly if they trust the public entity. Otherwise, they continue with the network formation model \mathcal{N} described in the previous section. An overview is given in Algorithm 4.

5.1 Fairness metric \mathcal{F}

Public entity P aims to maximize integration of new nodes in the network, and makes edge formation recommendations to achieve this goal. We formalize this in the form of fairness metric \mathcal{F} . To

Algorithm 4: Public entity P overview

```

Data:  $G = (V, E)$  the network of agents,  $\rho$  the public
entity's resource constraint
// SELECTION PROCESS
1  $V_{\text{ordered}} \leftarrow$  ordered nodes  $V$  according to iteration entered
the network ;
2  $V_{\text{selected}} \leftarrow V_{\text{ordered}}[:\rho]$  ; // select first  $\rho$ -many
// RECOMMENDATION PROCESS
3  $M \leftarrow \{\}$  ;
4 for  $u \in V_{\text{selected}}$  do
5    $v_{\text{best}} \leftarrow \operatorname{argmax}_{v \in V \setminus (\text{neighbors}(u) \cup u)} \mathcal{F}$  ; // choose
fairness-maximizing node
6    $u_{\text{trust}} \leftarrow$  True with probability  $\tau_u$  ;
7    $v_{\text{trust}} \leftarrow$  True with probability  $\tau_{v_{\text{best}}}$  ;
8   if  $u_{\text{trust}} \wedge v_{\text{trust}}$  then
9      $E \leftarrow E \cup (u, v)$  ; // if both nodes trust  $P$ ,
add edge
10  else
11     $\text{edgeFormation}(G, u)$ 
12 return  $G$ 

```

measure how well-integrated new nodes $v \in V_R$ are into the network, we choose to count how many triangles each node v is part of. There is empirical evidence to avoid low triangle counts – Bearman and Moody [7] find suicidal thoughts more likely in females who are socially isolated, and whose friends are not friends with each other. Our fairness metric is defined as follows, where V_R is the subset of agents that were not in the network at $t = 0$, and Δ_v is the number of triangles that v is part of. We normalize this value by the maximum possible number of triangles v could be part of: $\frac{(|V|-1)(|V|-2)}{2}$.

$$\mathcal{F} = \frac{2 \sum_{v \in V_R} \Delta_v}{|V_R|(|V|-1)(|V|-2)}$$

We recognize that this is not the only possible fairness metric for this application. Because our public entity P is not omniscient regarding the agents, we do not consider metrics that require calculation of utility or welfare. Under another formulation of this problem, calculating group fairness with respect to agent utility would be possible. It could even be done in a demographic-blind manner as described by Liu et al. [30]. We considered measures of access such as average path length, connectivity, and the methods described by Mehrotra et al. [31]. However, after exploring empirical studies [7, 12, 32] we decided that our agents need not have wide access to the network; the neighborhood is most important.

6 SIMULATION RESULTS

In this section, we answer our research questions by running simulations. We are not aware of an appropriate baseline algorithm to conduct comparisons, as this is a novel problem and setting. As a reminder, each agent v in the network has a different trust value τ_v , which represents the probability that they will take recommendation r_v from public entity P . The recommendations made by P intend to maximize fairness metric \mathcal{F} , which we have defined as

the proportion of possible triangles that new agents have formed. This fairness metric measures how well-integrated new arrivals are within the network. We prioritize integration because we study the application of refugees or other immigrants entering communities – without specific community knowledge they may be at a social disadvantage.

We simulate our network formation process (Section 4) in conjunction with the public entity model (Section 5). We run a simulation with 100 original agents V_O and 15 new agents V_R . Data from the American Immigration Council [15] informs these numbers; 13.7% of Americans are immigrants. Recall that each agent v also has a categorical attribute α_v . We draw α_v uniformly at random from two options for each agent v . These attributes are not the focus of our work, but are important in establishing agents' personal preferences for attributes in edge formation. The categorical attribute could represent a social interest or hobby in this work. As described in Section 4, new agents join one refugee per iteration. Once all new agents V_R have joined, we run the simulation for 50 additional iterations. We imagine each iteration to be at the time-scale of one week, so we collect network data after approximately one simulated year has passed.

6.1 RQ1: Trust and resource trade-off

In **RQ1** we ask: What is the trade-off between agent trust distribution T and the public entity's resource constraint ρ ? Can increasing the public entity's resources compensate for low agent trust as the public entity tries to improve \mathcal{F} every iteration?

To answer our research question, we vary ρ , the public entity's resource constraint, as well as the parameters of the Beta distribution for T . We test ρ values in $[0, 5, 10, 15, 20]$; we choose these values because we believe that it is reasonable for the public entity to have (approximately) sufficient resources only to help new arrivals. In figures, we present the public entity P 's resource constraint as being *proportional* to new agents V_R that the public entity is able to give recommendations to; we call these values x_ρ . We test trust Beta distributions in $[\beta(2, 20), \beta(2, 10), \beta(2, 5), \beta(2, 2), \beta(5, 2), \beta(10, 2), \beta(20, 2)]$. This results in the following mean trust values: $[0.09, 0.17, 0.29, 0.50, 0.71, 0.83, 0.91]$, which we call x_τ for clarity. Our trust values capture various high and low trust scenarios, and one moderate trust scenario. We show the results of our simulation in Figure 1. As expected, as we increase either ρ or mean trust, the public entity is able to achieve higher fairness \mathcal{F} with its recommendation algorithm \mathcal{A} . This is because increasing either of these values will increase the number of accepted recommendations by agents, therefore increasing the fairness metric \mathcal{F} .

Now, can increasing x_ρ make up for a low average agent trust x_τ ? We ask this question because in some cases, low agent trust levels in the public entity may be very difficult to change (though we explore some mechanisms to change agent trust in **RQ2** and **RQ3**). We can examine each row of Figure 1a to answer this question. In each row, average agent trust is held fixed, while the public entity's resources increase. When agent trust is low (e.g., the bottom three rows), increasing x_ρ has little influence on the fairness levels achieved (a 0.0006 unit increase, 0.0012 unit increase, and 0.0062 unit increase). For higher agent trust (though not the maximum), increasing public entity resources may indeed compensate for any lacking trust. Take

the second row from the top as an example, where average agent trust is 0.83. If the public entity is given enough resources, it can perform almost as well ($\approx 79\%$) as our maximal average agent trust.

In addition to an examination of the raw data, we run a multiple polynomial regression with independent variables of resource constraint x_ρ and average agent trust x_τ , to predict the fairness metric \mathcal{F} . We find the following relationship:

$$\mathcal{F} = 0.128x_\tau^2 - 0.001x_\rho^2 + 0.112x_\tau x_\rho - 0.118x_\tau - 0.023x_\rho + 0.024$$


with an r^2 value of 0.967. In Figure 1b we plot this function as a surface plot, with the fairness outcome \mathcal{F} represented both by the third axis and color. The fitted equation as well as the depiction in Figure 1b confirm that in low trust scenarios, increasing the resource constraint x_ρ has little effect. In fact, using the fitted equation, we can approximate the value of x_ρ required to achieve high fairness \mathcal{F} . Let's take the scenario where average agent trust is moderate: $x_\tau = 0.50$. What must x_ρ be to achieve maximal fairness \mathcal{F} (in this case, 0.14)? To answer this question, we solve: $0.14 = 0.128(0.50)^2 - 0.001x_\rho^2 + 0.112(0.50)x_\rho - 0.118(0.50) - 0.023x_\rho + 0.024$ for x_ρ . Our fitted function predicts that x_ρ must take on a value of 5.13 to achieve maximal fairness. Recall that in this section, x_ρ represents the *proportion* of new agents the public entity can help in a given iteration. Since in our implementation, the number of new agents $|V_R| = 15$, the public entity must be able to help $\rho = 15 * 5.13 \approx 77$ agents per iteration to achieve maximal trust. This is approaching the size of the entire network, and not a feasible resource constraint for our public entity. We also note that our function does not have any real roots when x_τ takes on lower values (0.29 and 0.17) — the

prediction is that it is *impossible* to achieve maximal fairness for low trust values. Thus we can confidently answer: *no*, increasing the public entity's resources ρ *cannot* compensate for low agent trust T .

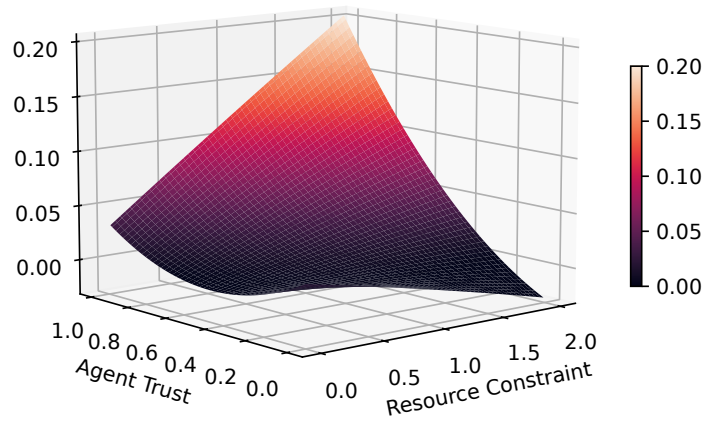
6.2 RQ2: Targeted advertising campaign

In RQ2 we ask: If an agent's trust level τ_v is known by the public entity, can P target agents with low trust via advertising to increase trust and fairness \mathcal{F} ? How much money would P spend on this campaign?

To test this research question, we run simulations under the same parameters described previously. These simulations, however, require that the public entity P divert some proportion of its resources ρ to target low-trust agents. We call this proportion q . We imagine that the public entity P has close ties to the community and interacts often with individuals — through these interactions the public entity can gain information regarding who the low-trust individuals are. Then, public entity P can target $(q * \rho)$ -many of these low-trust agents each iteration, while still helping $((1 - q) * \rho)$ -many agents via recommendations. The targeted campaign proceeds as follows. Each iteration, public entity P allocates q proportion of its resources to the campaign. The public entity chooses the $(q * \rho)$ -many agents with the lowest trust to target. Note that it is important for all agents to trust the public entity — even those who are not being given recommendations. If P recommends that agent v form an edge with u , agent u must also have trust in the public entity for this edge to form. For each agent v that has been selected, the public entity calculates the maximal trust improvement $(1 - \tau_v)$. The public entity



(a) Fairness \mathcal{F}



(b) Regression on fairness \mathcal{F}

Figure 1: Fairness \mathcal{F} as we vary the proportion of new agents the public entity can help, x_ρ , and average trust x_τ . A combination of high resources average trust lead to the best fairness outcomes (a). It is difficult to achieve increased fairness when trust is low simply by increasing resources. We fit a multiple polynomial regression to the data in (a), which we show as a surface plot in (b). The fitted equation is given by: $\mathcal{F} = 0.128x_\tau^2 - 0.001x_\rho^2 + 0.112x_\tau x_\rho - 0.118x_\tau - 0.023x_\rho + 0.024$, and the r^2 value is 0.967. This regression indicates that it is infeasible to increase resources enough to compensate for low trust.




Figure 2: We test advertising interventions with $x_\rho = 1$ and various values of q (the proportion of resources diverted to advertising). In (a) we see how these interventions affect fairness \mathcal{F} ; in cases of low agent trust (below 0.4), most interventions outperform the null case, due to successful improvement of average agent trust (b). The amount of money spent by the public entity (c) does not differ much for varying q values, but decreases as x_τ increases. The black dotted line in each figure shows the original metric.

P will spend money proportional to the trust improvement ($\$1 - \tau_v$) to advertise to agent v . We make an assumption that the money put into advertising directly translates to a successful campaign. After the advertising campaign, each targeted agent’s new trust value is as follows: $\tau_v \rightarrow \tau_v + c(1 - \tau_v)$, where c is the proportion of the agent’s “trust gap” that is restored by the campaign. For our model, we set $c = 0.71$. This is justified by empirical data [17]; we describe these experiments in more detail in the next section.

We show results of these targeted advertising interventions in Figure 2. We fix $x_\rho = 1$ (or, $\rho = 15$). As q increases, more resources are diverted from ρ to be used in the advertising campaign. Figure 2a shows us that for low average trust ($x_\tau < 0.4$), most interventions perform better than no interventions at all. The exception is the $q = 5/5 = 1$ case. This is to be expected, as the public entity P is diverting all of its resources from recommendations. While agents in this case might be very trusting of P (see Figure 2b), that trust is not put to good use, as no recommendations are being made. For higher average trust levels ($x_\tau \geq 0.4$), interventions with lower q values perform better, until about $x_\tau = 0.84$, where the no intervention case dominates. From this figure we can see that lower values of q perform better — it is important to keep making recommendations to agents in order to achieve fairness \mathcal{F} . The price of the intervention increases with q ; however these values all track quite closely (Figure 2c). As expected, the amount spent on the intervention is quite low when average trust x_τ is high. However, as mentioned previously, the fairness payoff diminishes as trust increases.

6.3 RQ3: Transparent public entity

In RQ3 we ask: Consider a schema where the public entity P announces the effects of its interventions on fairness. If an agent’s trust level τ_v is affected by this announcement, can transparency from P lead to improved fairness \mathcal{F} ?

We run simulations under the same parameters. To answer this research question, we introduce a mechanism where public entity

P announces how the network has changed under its recommendation system, after which the agents update their trust. We consider two mechanisms — a global announcement and update, and local.

Globally transparent entity: Under the global mechanism, public entity P announces at the beginning of each iteration how fairness \mathcal{F} has changed since the prior iteration. Call this change in fairness $\delta_{\mathcal{F}}$. All agents $v \in V$ then update their trust values τ_v according to the announced value $\delta_{\mathcal{F}}$. Note that $\delta_{\mathcal{F}}$ can be negative if the public entity does not succeed, thus agent trust can decrease. Each iteration, each agent v ’s trust value is updated as follows: $\tau_v \rightarrow \tau_v + c_g \delta_{\mathcal{F}}$, where c_g is the proportion of trust repaired by a global announcement. This linear model is derived from experiments conducted by Desmet [17]. Desmet studied how a monetary reward might repair trust between two agents if one was harmed by the other; and found a linear relationship. Specifically, if the harm was perceived as ambiguous (rather than intentional), then partial compensation for this harm would increase trust by a factor of 0.71. Since we consider this global announcement to represent an impersonal public entity, and since increasing fairness globally will only partially (rather than completely) compensate agents for low trust, we use $c_g = 0.71$ in our simulations.

Locally transparent entity: Under the local mechanism, the public entity P makes a personalized announcement to each node $v \in V$ at the beginning of each iteration. In practice, this announcement might be a personalized monthly statement, which we do not foresee being resource-intensive. Public entity P announces how fairness metric \mathcal{F} has changed since the beginning of the prior iteration in each node’s neighborhood. Call this change in fairness in node v ’s neighborhood $\delta_{\mathcal{F}_v}$. To calculate this value, we use the subgraph induced by the neighborhood of each node (including that node) instead of the entire graph G . If the subgraph contains fewer than three nodes, $\mathcal{F} = 0$. Then, each node v updates their trust value as follows: $\tau_v \rightarrow \tau_v + c_l \delta_{\mathcal{F}_v}$, where c_l is the proportion of trust repaired by a local announcement. We use the same experimental




Figure 3: Fairness \mathcal{F} for different values of public entity P 's resource constraint x_ρ and mean agent trust x_ρ , for a global (a) and local (b) announcement mechanism. The local mechanism more successfully increases fairness in low trust scenarios. We show how mean agent trust is affected across these values for the global (c) and local (d) mechanisms. Each row label gives the original mean agent trust. Notice that resource constraint x_ρ has a small positive effect on this outcome in (c) and a small negative effect in (d).

results [17] to justify this model; if the harm was perceived as intentional, then a partial compensation for this harm would only increase trust by a factor of 0.48. We consider the local announcement to be personal, and since increasing fairness will still only partially compensate agents, we use $c_l = 0.48$ in our simulations.

Figure 3 shows the results from the implementation of these two mechanisms. First, we make a comparison to Figure 1a. The global

announcement mechanism (Figure 3a) makes a small improvement to fairness \mathcal{F} , concentrated in high trust scenarios. Figure 3c confirms small improvements in trust overall. Why might this be the case? The global announcement is at a very large scale, and impersonal. Because improvement to \mathcal{F} happens slowly in small increments, the value $\delta_{\mathcal{F}}$ is quite small at each iteration. We would not expect to see a large jump in fairness or trust unless the public

entity caused very large improvements to the network in one iteration. In comparison, the local mechanism (Figure 3b) *does* cause larger fairness improvements in low trust scenarios. If we look at the case where $x_\tau = .50$, and $x_\rho = 1.33$, fairness \mathcal{F} increased from 0.031 to 0.06; a 94% improvement. Looking at the mean trust x_τ under the local mechanism (Figure 3d), the improvements here were much greater in low trust scenarios. Again, this is due to the scale of the local fairness improvements; in a small local neighborhood, the improvement in \mathcal{F} is on a much larger scale. This is what we want in our model, as we assume agents care more about their local neighborhoods than the entire community. However, this comes with a trade-off that any decreases in fairness will also have a large impact on trust. We hypothesize that this is leading to less effective trust improvements in the high x_ρ scenarios.

Our results give theoretical insights regarding refugee resettlement. First; trust is essential — it can drastically impact whether individuals follow advice from authority figures. Second, one’s integration in a community *can* change with more information or resources. Public entities’ involvement in the lives of migrants in this way could prove extremely useful. And third, if the public entity spends resources (time, energy, money) on low trust individuals, they can be swayed. This increased trust can have positive impacts on individuals and communities.

7 DISCUSSION

Network formation model: We adapt the network formation model from Christakis et al. [14] as our basis for agent behavior. It dictates the desires of community members, and determines how they will act if not given a recommendation by algorithm \mathcal{A} ; with such small resource constraints ρ , most agents are not. We use a model from prior work because we desire a model that has already been proven to be ecologically valid. It is not the key to our contributions, but it affects our results nonetheless. If we had used a different network formation model, a number of things could have changed. First, consider the alignment in desires of the agents and the public entity P . The public entity wants to increase fairness \mathcal{F} . We define this to be the average number of triangles each new agent is able to form (normalized by the total possible number of triangles) — see Section 5.1. Each agent also has a utility function f (see Section 4), which is a linear combination of various factors, including whether or not the proposed edge will form a triangle. In our implementation, the coefficient for this factor is positive, meaning that both the agents and the public entity are incentivized to form triangles. In this work, we did not explore the scenario where agents and the public entity are expressly opposed, but it would make for interesting future work. Second, agents in this model are not resource bounded — they can form as many edges as they like. This means that the public entity need not consider agents’ current degrees when making edge formation recommendations. This addition, while more ecologically valid, could prove to be a more nuanced optimization problem, better suited for future work.

Interventions’ effects on trust: We assume when investigating RQ2 and RQ3 that 1) advertising campaigns and 2) announced increased fairness *positively* affect trust. However, even well-funded advertising campaigns can fail. In fact, a poor advertising

campaign (where agents feel overly-targeted, perhaps) may actually result in a loss of trust in the public entity. Modeling of these campaigns is outside of the scope of this work, but would make for interesting future experimental or survey work. We also assume that the announced increase in fairness has a positive effect on agent trust. However, it may be the case that some agent has entirely selfish desires, and gains nothing from learning about increased fairness in the community. In this case, this agent might require a personalized statement describing how the public entity helped them specifically. This would be cost intensive and difficult to do on a large scale.

Public entities with ill-intent: In this work, we assume an altruistic public entity aims to increase fairness in the network. We do not intend for our answer to RQ2 (targeted advertising campaign) to be used to achieve some goal of an entity other than the good of the community. We ask and investigate RQ3 as a safeguard against public entities with ill-intent. We encourage transparency (especially locally) so that community members can know exactly what kind of impact the public entity is having on their local communities. Agents can then make their own decisions about whether to trust and take recommendations from the public entity. We also believe that skepticism will safeguard agents against a public entity that wishes to harm them. While in our model we assume that a trusting agent will blindly take the recommendations of a public entity, we know that real agents will consider this decision more carefully. A future modeling problem might focus on more skeptical community members.

8 LIMITATIONS AND FUTURE WORK

Agent-based modeling: Though we draw inspiration heavily from prior empirical work, agent-based modeling can never tell us the whole story. We make many simplifying assumptions regarding time-scale, work done by the public entity P , and attitudes and behaviors of agents. This work does, however, give theoretical evidence for future work. We find that two interventions are successful, and we hope that future experimental work could explore design factors such as phrasing of the announcement or the type of advertisement.

Trust as a social phenomenon: In this work, we make a simplifying assumption that an individual’s trust in an entity is personal and fixed. We know, however, that information and opinions cascade quickly, and the opinion of one’s neighbor might strongly influence one’s own opinions. In particular, if the public entity gives a particularly good (or bad) recommendation to a refugee, they might be inclined to tell other community members about it. This information (like entity transparency) might have a large affect on trust. A model that incorporates this aspect of human behavior would make for very interesting future work.

Alternate recommendation frameworks: We fix recommendation algorithm \mathcal{A} , and make no claims that we have found the *ideal* algorithm. However, future work could take on this optimization problem. In addition, we assume that new agents in the network desire and seek help from the public entity. However, we might instead model agents who are hesitant to ask or don’t have the resources to do so. In another framework, the public entity P might use its resources to recommend one edge multiple

times in one iteration, specifically in scenarios of low trust, if it will cause a large increase in fairness.

9 CONCLUSION

In this work we introduce a framework for refugee integration into existing communities. We present an altruistic public entity P that operates within a community of agents V . The public entity gives recommendations to agents who engage in a strategic network formation process. These recommendations come in the form of edge formation suggestions which aim to increase integration of refugees. An agent v will choose to take a recommendation with probability τ_v , the trust agent v has in the entity. We show that increasing the public entity's resource allocation will *not* compensate for low trust in the community; with low trust, P will not be able to achieve its fairness goal \mathcal{F} . We introduce two trust-increasing interventions, 1) a targeted advertising campaign, and 2) an announcement from a transparent public entity. We find that diverting a fraction (20%) of the public entity's resources to a targeted advertising campaign can increase trust and fairness in the community, especially in low community trust scenarios (under $x_r = 0.40$). We also find that personalized announcements regarding fairness in local neighborhoods are more effective in increasing trust and fairness in low trust scenarios than global announcements. Local announcements can almost double fairness metric \mathcal{F} in some cases. Our work underscores the importance of community trust in integration; low trust cannot be compensated for with greater resources. Importantly, transparency of the public entity requires no extra resource expenditure, but can increase fairness. This work provides theoretical support for these trust-increasing interventions, which we showed via simulation lead to increased integration of refugees in communities.

ACKNOWLEDGMENTS

We would like to thank Aditya Karan, Ali Zaidi, and the anonymous reviewers for their many useful comments and suggestions.

REFERENCES

- [1] Narges Ahani, Paul Gözl, Ariel D Procaccia, Alexander Teytelboym, and Andrew C Trapp. 2023. Dynamic placement in refugee resettlement. *Operations Research* (2023).
- [2] Réka Albert and Albert-László Barabási. 2002. Statistical mechanics of complex networks. *Reviews of modern physics* 74 (2002), 47.
- [3] Junaid Ali, Mahmoudreza Babaei, Abhijnan Chakraborty, Baharan Mirzasoileman, Krishna Gummadi, and Adish Singla. 2021. On the fairness of time-critical influence maximization in social networks. *IEEE Transactions on Knowledge and Data Engineering* (2021).
- [4] James Andreoni. 1989. Giving with impure altruism: Applications to charity and Ricardian equivalence. *Journal of political Economy* 97, 6 (1989), 1447–1458.
- [5] James Andreoni. 1990. Impure altruism and donations to public goods: A theory of warm-glow giving. *The economic journal* 100, 401 (1990), 464–477.
- [6] James Atwood, Hansa Srinivasan, Yoni Halpern, and David Sculley. 2019. Fair treatment allocations in social networks. *arXiv preprint arXiv:1911.05489* (2019).
- [7] Peter S Bearman and James Moody. 2004. Suicide and friendships among American adolescents. *American journal of public health* 94, 1 (2004), 89–95.
- [8] Morton Beiser. 2006. Longitudinal research to promote effective refugee resettlement. *Transcultural psychiatry* 43, 1 (2006), 56–71.
- [9] Morton Beiser. 2009. Resettling refugees and safeguarding their mental health: Lessons learned from the Canadian Refugee Resettlement Project. *Transcultural psychiatry* 46, 4 (2009), 539–583.
- [10] Renaud Bourlès, Yann Bramoullé, and Eduardo Perez-Richet. 2017. Altruism in networks. *Econometrica* 85, 2 (2017), 675–689.
- [11] Anastasia Brown and Todd Scribner. 2014. Unfulfilled promises, future possibilities: The refugee resettlement system in the United States. *Journal on Migration and Human Security* 2, 2 (2014), 101–120.
- [12] Raffaella Calati, Chiara Ferrari, Marie Brittner, Osmano Oasi, Emilie Olié, André F Carvalho, and Philippe Courtet. 2019. Suicidal thoughts and behaviors and social isolation: A narrative review of the literature. *Journal of affective disorders* 245 (2019), 653–667.
- [13] Hsiu-Ling Chen, Hsueh-Liang Fan, and Chin-Chung Tsai. 2014. The role of community trust and altruism in knowledge sharing: An investigation of a virtual community of teacher professionals. *Journal of Educational Technology & Society* 17, 3 (2014), 168–179.
- [14] Nicholas Christakis, James Fowler, Guido W Imbens, and Karthik Kalyanaraman. 2020. An empirical model for strategic network formation. In *The Econometric Analysis of Network Data*. Elsevier, 123–148.
- [15] American Immigration Council. 2021. Immigrants in the United States. <https://map.americanimmigrationcouncil.org/locations/national/#:~:text=Immigrants%20in,United%20States%20of%20America&text=The%20United%20States%20was%20built,of%20whom%20are%20naturalized%20citizens>.
- [16] David Delacrétaz, Scott Duke Kominers, and Alexander Teytelboym. 2016. Refugee resettlement. *University of Oxford Department of Economics Working Paper* (2016).
- [17] Pieter Desmet. 2011. *In money we trust? Trust repair and the psychology of financial compensations*. Number EPS-2011-232-ORG.
- [18] David DeSteno. 2015. Compassion and altruism: How our minds determine who is worthy of help. *Current opinion in behavioral sciences* 3 (2015), 80–83.
- [19] Immacolata Di Napoli, Pasquale Dolce, and Caterina Arcidiacono. 2019. Community trust: A social indicator related to community engagement. *Social Indicators Research* 145 (2019), 551–579.
- [20] Paul Erdős and Alfréd Rényi. 1960. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci* 5 (1960), 17–60.
- [21] Victoria M Esses, Leah K Hamilton, and Danielle Gaucher. 2017. The global refugee crisis: Empirical evidence and policy implications for improving public attitudes and facilitating refugee resettlement. *Social issues and policy review* 11, 1 (2017), 78–123.
- [22] Ernst Fehr and Urs Fischbacher. 2003. The nature of human altruism. *Nature* 425, 6960 (2003), 785–791.
- [23] Daniel C Feiler, Leigh P Tost, and Adam M Grant. 2012. Mixed reasons, missed givings: The costs of blending egoistic and altruistic reasons in donation requests. *Journal of Experimental Social Psychology* 48, 6 (2012), 1322–1328.
- [24] Paul Gözl and Ariel D Procaccia. 2019. Migration as submodular optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 549–556.
- [25] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. 1983. Stochastic blockmodels: First steps. *Social networks* 5, 2 (1983), 109–137.
- [26] Jon M Jachimowicz, Salah Chafik, Sabeth Munrat, Jaideep C Prabhu, and Elke U Weber. 2017. Community trust reduces myopic decisions of low-income individuals. *Proceedings of the National Academy of Sciences* 114, 21 (2017), 5401–5406.
- [27] Zeinab S Jalali, Qilan Chen, Shwetha M Srikanta, Weixiang Wang, Myunghwan Kim, Hema Raghavan, and Sucheta Soundarajan. 2023. Fairness of Information Flow in Social Networks. *ACM Transactions on Knowledge Discovery from Data* 17, 6 (2023), 1–26.
- [28] Zeinab S Jalali, Weixiang Wang, Myunghwan Kim, Hema Raghavan, and Sucheta Soundarajan. 2020. On the information unfairness of social networks. In *Proceedings of the 2020 SIAM International Conference on Data Mining*. SIAM, 613–521.
- [29] Stephen Judd, Michael Kearns, and Yevgeniy Vorobeychik. 2011. Behavioral conflict and fairness in social networks. In *Internet and Network Economics: 7th International Workshop, WINE 2011, Singapore, December 11-14, 2011. Proceedings* 7. Springer, 242–253.
- [30] David Liu, Virginie Do, Nicolas Usunier, and Maximilian Nickel. 2023. Group fairness without demographics using social networks. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1432–1449.
- [31] Anay Mehrotra, Jeff Sachs, and L Elisa Celis. 2022. Revisiting Group Fairness Metrics: The Effect of Networks. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–29.
- [32] Chloé Motillon-Toudic, Michel Walter, Monique Séguin, Jean-Daniel Carrier, Sofian Berrouguet, and Christophe Lemey. 2022. Social isolation and suicide risk: Literature review and perspectives. *European psychiatry* (2022), 1–42.
- [33] Stephanie J Nawyn. 2006. Faith, ethnicity, and culture in refugee resettlement. *American Behavioral Scientist* 49, 11 (2006), 1509–1527.
- [34] Bich-Ngan T Nguyen, Phuong NH Pham, Van-Vang Le, and Václav Šnášel. 2022. Influence maximization under fairness budget distribution in online social networks. *Mathematics* 10, 22 (2022), 4185.
- [35] Office of Refugee Resettlement. 2023. About the Refugee Program. <https://www.acf.hhs.gov/orr/about-refugee-program>
- [36] Office of Refugee Resettlement. 2023. Resettlement Services. <https://www.acf.hhs.gov/orr/programs/refugees>
- [37] Douglas Paton. 2007. Preparing for natural hazards: the role of community trust. *Disaster Prevention and Management: An International Journal* 16, 3 (2007), 370–379.
- [38] Behnam Razaghi, Mehdy Roayaei, and Nasrollah Moghadam Charkari. 2022. On the Group-Fairness-Aware Influence Maximization in Social Networks. *IEEE*

- Transactions on Computational Social Systems* (2022).
- [39] Fernando P Santos, Jorge M Pacheco, Ana Paiva, and Francisco C Santos. 2017. Structural power and the evolution of collective fairness in social networks. *PLoS one* 12, 4 (2017), e0175687.
- [40] Akрати Saxena, George Fletcher, and Mykola Pechenizkiy. 2022. Fairsna: Algorithmic fairness in social network analysis. *arXiv preprint arXiv:2209.01678* (2022).
- [41] Jeffrey R Stevens and Marc D Hauser. 2004. Why be nice? Psychological constraints on the evolution of cooperation. *Trends in cognitive sciences* 8, 2 (2004), 60–65.
- [42] Ana-Andreea Stoica and Augustin Chaintreau. 2019. Fairness in social influence maximization. In *Companion Proceedings of The 2019 World Wide Web Conference*. 569–574.
- [43] Joanne Van Selm. 2014. Refugee resettlement. (2014).
- [44] Andreas Vårheim. 2014. Trust and the role of the public library in the integration of refugees: The case of a Northern Norwegian city. *Journal of Librarianship and Information Science* 46, 1 (2014), 62–69.
- [45] Andreas Vårheim. 2014. Trust in libraries and trust in most people: Social capital creation in the public library. *The Library Quarterly* 84, 3 (2014), 258–277.
- [46] Andreas Vårheim, Sven Steinmo, and Eisaku Ide. 2008. Do libraries matter? Public libraries and the creation of social capital. *Journal of documentation* 64, 6 (2008), 877–892.
- [47] Susanne Wallman Lundåsen and Dag Wollebæk. 2013. Diversity and community trust in Swedish local communities. *Journal of Elections, Public Opinion & Parties* 23, 3 (2013), 299–321.
- [48] Duncan Watts and Steven Strogatz. 1998. Collective dynamics of ‘small-world’ networks. *Nature* 393 (1998), 440–442.
- [49] Monica Webb Hooper, Charlene Mitchell, Vanessa J Marshall, Chesley Cheatham, Kristina Austin, Kimberly Sanders, Smitha Krishnamurthi, and Lena L Grafton. 2019. Understanding multilevel factors related to urban community trust in healthcare and research. *International journal of environmental research and public health* 16, 18 (2019), 3280.
- [50] Norman L Zucker. 1983. Refugee resettlement in the United States: Policy and problems. *The Annals of the American Academy of Political and Social Science* 467, 1 (1983), 172–186.