Taylor & Francis
Taylor & Francis Group

INEQUALITY AND NETWORK FORMATION GAMES

Samuel D. Johnson and Raissa M. D'Souza

Department of Computer Science, University of California, Davis, Davis, California, USA

Abstract *This article addresses the matter of inequality in network formation games. We employ a quantity that we are calling the* Nash Inequality Ratio *(NIR), defined as the maximal ratio between the highest and lowest costs incurred to individual agents in a Nash equilibrium strategy, to characterize the extent to which inequality is possible in equilibrium. We give tight upper bounds on the NIR for the network formation games of Fabrikant et al. [14] and Ehsani et al. [13]. With respect to the relationship between equality and social efficiency, we show that, contrary to common expectations, efficiency does not necessarily come at the expense of increased inequality.*

1. INTRODUCTION

Noncooperative game theory uses the concept of equilibria to capture the idea that, in a competitive world, rational agents will maneuver themselves to a fixed point from which no further maneuvering will yield additional benefits (e.g., a lower cost). The most ubiquitous equilibrium concept is the *Nash equilibrium*, which is satisfied when no individual agent can achieve a lower cost by changing its strategy given that the strategies of every other agent remain unchanged. In a Nash equilibrium there can exist a disparity between the costs incurred by individual agents, with the "more fortunate" agents subjected to lower costs than the "less fortunate" ones. In this manuscript, we investigate the *Nash Inequality Ratio* (NIR), defined as the maximum value of the ratio between the highest and lowest costs found within a single Nash equilibrium, to determine the extent to which cost disparity can arise between pairs of agents in an equilibrium outcome.

Recent years have witnessed widespread, growing concerns over economic inequality (see, for example, [9, 31, 29]). Yet, within the the algorithmic game theory community—a field of study sitting at the intersection of computer science, economics, and game theory—inequality has received very little attention, as evidenced by the nonexistence of a standard vernacular with which to discuss bounds of the sort captured by the NIR. The NIR expresses, in a natural way, a fundamental property of a strategic scenario, akin to the well-known metrics used to characterize other qualities of equilibria such as the Price of Anarchy (PoA) [22, 23] and the Price of Stability (PoS) [3].

Address correspondence to Samuel D. Johnson, Department of Computer Science, University of California, Davis, One Shields Avenue, Davis, CA 95616, USA. E-mail: samjohnson@ucdavis.edu

In this manuscript we analyze inequality in network formation games. Network formation games model the formation of network structures by and between a collection of strategic, self-interested individuals or *agents*. In these games, connectivity is deemed to be desirable though costly, and it is up to the agents to reconcile the gains they achieve through additional connectivity (e.g., access to information, the ability to communicate and coordinate, etc.) with the costs or resource limitations that limit the number of links that they can afford to create.

We focus on two network formation games: the Undirected Connections (UC) game [14] and the Undirected Bounded Budget Connections (UBBC) game [13]. In both games, there is a set $N = \{1, \ldots, n\}$ of strategic agents. Each agent $i \in N$ selects a linking strategy $s_i \subseteq N \setminus \{i\}$ that identifies a subset of other agents to which they will build links. The joint strategy $s = (s_1, \ldots, s_n)$ induces an undirected network $G_s = (N, E_s)$ in which the agents themselves are vertices and the edge set E_s is the union of every agent's linking strategies. In both games, agents incur a *usage cost* that is defined for each $i \in N$ to be the sum of distances between i and every other $j \in N \setminus \{i\}$ in G_s. The two games differ in how the cost of building edges is accounted for: in the UC game, each agent i incurs an additional *construction cost*, defined to be α times the number of edges that i builds; and in the UBBC game, each agent i is endowed with an *edge budget*, $k_i > 0$, determining the maximum number of edges that i can build.

Our first results are presented in Section 3 where we establish tight upper bounds on the NIR for these two network formation games. For the UC game, we establish bounds parameterized on the cost of building links, α, which is assumed to be a constant. In particular, when links are inexpensive ($\alpha < 1$), we find that the NIR is at most $1 + \alpha < 2$; and when links are expensive ($\alpha \geq 1$), the NIR is asymptotically bounded (as the number of agents $n \to \infty$) by $\max\{2, (1 + \alpha)/2\}$. Then, for the UBBC game, we prove that the NIR is bounded by the constant 2. We show that this bound is asymptotically tight for every positive budget allocation—including those that endow different budgets to different agents.

With these results established, in Section 4 we examine the relationships between efficiency and (in)equality. Here we focus on Nash equilibrium strategies that are also *efficient* (i.e., minimize the social cost). In both games we find that this relationship is largely dependent on the availability of resources (edges). In particular, we find that when resources are scarce (that is, when edges are expensive in the UC model and when budgets are small in the UBBC model), the two games behave quite differently from each other. In the scarce regime for the UBBC game, no efficient Nash equilibrium strategies achieve cost equality, and some actually maximize the inequality ratio (achieving the NIR upper bound established in Section 3). To the contrary, in the scarce regime for the UC game, there are some efficient Nash equilibrium strategies with egalitarian costs and others with maximal inequality. These results demonstrate that the relationship between equality, efficiency, and equilibrium is entirely model specific and varies across network formation games.

The research agenda initiated in this manuscript—using the NIR to characterize the relationship between equality, efficiency, and equilibrium—opens an avenue for future research that involves the bounding of these quantities in different games. Further, we hope that this work stimulates interest in examining questions of inequality and equilibrium more broadly.

Section 5 concludes with closing remarks. Finally, to facilitate the flow of the main ideas, tangential proofs are deferred to the Appendix.

2. PRELIMINARIES

In this section, we formally define the Nash Inequality Ratio (Section 2.1) and the network formation games (Section 2.2) on which the remainder of the article is focused.

2.1. Nash Inequality Ratio

Studies of economic inequality have spawned a number of metrics that aim to quantify the level of inequality in a given system, the most well-known being the Gini coefficient [15] used to quantify the level of inequality in the distribution of utility (wealth) across a population of individuals. The related topic of *fairness* has also received considerable attention in the game theory literature, particularly with respect to mechanism design [33, 10, 21] in which the goal is to develop allocation mechanisms that achieve various notions of equity among the parties involved. Roughgarden [30] proposed a metric quantifying *unfairness* in the context of noncooperative routing games as the maximum ratio between an agent's cost (i.e., latency of a flow path) in a socially optimal outcome and that of a Nash equilibrium outcome.

In this manuscript, we put forward the NIR as a simple metric that bounds, in a natural way, the extent that inequality between agents is supported in Nash equilibrium outcomes.

Consider a game Γ involving a set $N = \{1, \ldots, n\}$ of players/agents where, for each agent $i \in N$, S_i specifies i's strategy space. A joint strategy $s = (s_1, \ldots, s_n) \in S_1 \times \cdots \times S_n = S$ yields an outcome, and each agent i incurs a cost $c_i(s)$ that is a function of the joint strategy.

Given a joint strategy s, let $\mu = \mathrm{argmin}_{i \in N}\, c_i(s)$ denote an agent that incurs a minimal cost, and $\chi = \mathrm{argmax}_{i \in N}\, c_i(s)$ denote an agent that incurs the highest cost. The *inequality ratio* (IR) for the joint strategy s is defined as:

$$\mathrm{IR}(s) = \frac{c_\chi(s)}{c_\mu(s)} = \frac{\max_{i \in N} c_i(s)}{\min_{j \in N} c_j(s)}.$$

That is, the IR of a strategy s is the maximal cost ratio between any pair of agents.

Recall that a joint strategy s is a Nash equilibrium if, for every agent $i \in N$ and every $s_i' \in S_i$, we have

$$c_i(s_i, s_{-i}) \le c_i(s_i', s_{-i}),$$

where s_{-i} is shorthand for the joint strategy of every agent except i. In a Nash equilibrium, no agent stands to lower its cost through a unilateral strategy update. Let $S_{NE} \subseteq S$ denote the set of all Nash equilibrium. The NIR is defined to be the greatest IR among all Nash equilibrium strategies:

$$\mathrm{NIR}(\Gamma) = \max_{s \in S_{NE}} \mathrm{IR}(s). \tag{2.1}$$

Note that the NIR is kind of a "max-max" quantity, in that the NIR of a game is the maximum inequality ratio among all Nash equilibrium strategies for that game, and each inequality ratio itself is a maximal quantity with respect to a particular strategy.

Finally, by definition we have $\text{IR}(s) \geq 1$ for every strategy $s \in S$. A strategy s is called *egalitarian* if $\text{IR}(s) = 1$.

2.2. Network Formation Games

Network formation games model the creation of networks by rational and self-interested agents strategically building edges between one another. In this section, we describe the two specific network formation games on which we base the inequality analysis presented in this article: the Undirected Connections game (Section 2.2.1) and the Undirected Bounded Budget Connections game (Section 2.2.2). For a comprehensive introduction to network formation games the reader is referred to the surveys of [18, 19, 16, 32].

Questions regarding inequality in network formation games have received very little attention in the literature. To the best of our knowledge, the only study that looks explicitly at inequality in noncooperative network formation games is found in [17], who analyze the effect of local splillovers in pairwise stable networks. The authors found that agents with more connections and larger neighborhoods earn higher utilities than those with fewer connections and smaller neighborhoods. The results presented in [17] show that inequality can arise in equilibrium outcomes, but the authors do not go on to quantify the extent of the inequality.

2.2.1. The Undirected Connections Game. The UC network formation game[1] introduced by Fabrikant et al. [14] was the first network formation game to appear in the algorithmic game theory literature.[2] In addition to introducing the UC model and establishing some basic properties of efficient and Nash equilibrium outcomes, Fabrikant et al. establish bounds on the PoA, which have subsequently been tightened (e.g., [1, 11, 27, 25]).

The UC model, as defined by [14], is specified by a set $N = \{1, \ldots, n\}$ of strategic agents and a parameter $\alpha > 0$, a constant, that determines the cost of building a single edge. The strategy space for an individual agent $i \in N$, denoted S_i, consists of all possible subsets of other agents to which i can build a direct connection; i.e., $S_i \subseteq \mathcal{P}(N \setminus \{i\})$.[3] A (pure) strategy for agent i, denoted $s_i \in S_i$, is a specific subset of other agents with which i wishes to establish links. A joint strategy $s = (s_1, s_2, \ldots, s_n)$, representing the strategy selections of every agent in N, induces an undirected network $G_s = (N, E_s)$, wherein the agents themselves are present as vertices and the edge set is defined to be $E_s = \{\{i, j\} : j \in s_i\}$. Because an undirected edge $\{i, j\}$ is present in G_s if either $j \in s_i$ or $i \in s_j$, edge formation in the UC model is said to be *unilateral*.

Each agent i incurs a cost that is a function of both its own strategy, s_i, and the joint strategy of every other agent, s_{-i}. This cost for agent i consists of both a *usage cost*, $d_i(s)$, and a *creation cost*, $b_i(s)$. The usage cost is defined to be the sum of distances between i and every other agent j; $d_i(s) = \sum_{j \in N} \ell_{G_s}(i, j)$ where $\ell_{G_s}(i, j)$ denotes the length of the

---

[1]The name Undirected Connections is our own choosing in order to reinforce the nature of the game.

[2]We note that earlier work on network formation games appeared in the economics literature—cf. [4, 20, 5].

[3]We use the notation $\mathcal{P}(X)$ to denote the *power set* of a set X; i.e., the set of all subsets of X.

shortest path between nodes i and j in the graph G_s (or ∞ if no such path exists).[4] The creation cost is defined to be linear in the number of edges i contributes to the network's construction; specifically, $b_i(s) = \alpha \cdot |s_i|$ where $|s_i|$ conveys the number of edges that i builds in the network and $\alpha \geq 0$ is a constant, specified as a game parameter. Hence, the cost to an agent $i \in N$ given the joint strategy profile $s = (s_i, s_{-i})$ is

$$c_i(s) = b_i(s) + d_i(s) = \alpha|s_i| + \sum_{j \in N} \ell_{G_s}(i, j). \qquad (2.2)$$

The *social cost* is defined as the sum of the agents' individual costs;

$$C(s) = \sum_{i \in N} c_i(s). \qquad (2.3)$$

A strategy profile $s \in S$ that minimizes (2.3) is called *efficient*.

Fabrikant et al. identify the topologies of efficient and Nash equilibrium strategies for three different regimes of α, which we summarize in Proposition 2.1. These results will play prominently in our analysis of the UC game.

Proposition 2.1. [14] *Efficient and Nash equilibrium outcomes for the UC network formation game are as follows:*

1. *When $\alpha < 1$, the complete graph is both efficient and the only Nash equilibrium.*[5]
2. *When $1 \leq \alpha < 2$, the complete graph is efficient but the star is the only Nash equilibrium.*[6]
3. *When $2 \leq \alpha$, the star is efficient and a Nash equilibrium, although there are other Nash equilibrium outcomes as well.*

2.2.2. The Undirected Bounded Budget Connections Game.

The UBBC network formation game was introduced by Ehsani et al. [13] as an undirected variant of the Bounded Budget Connections (BBC) game of [24]. Ehsani et al. establish upper bounds on the PoA for the UBBC game, building on the techniques developed by [2] for a related network formation game.

As with the UC game, the UBBC game involves a set $N = \{1, \ldots, n\}$ of strategic agents building edges between one another. However, in the UBBC model, each agent $i \in N$ is endowed with a *budget* $k_i > 0$ that determines the maximum number of edges the agent can create. The budgets $k_1, k_2, \ldots, k_n$ are specified exogenously and, in general, need not be identical. We refer to the special case when all edge budgets are identical (i.e., $\forall i \in N, k_i = k$) as *uniform* instances of the UBBC game.

Edge formation is unilateral, so a joint strategy $s = (s_1, \ldots, s_n)$ induces a network $G_s = (N, E_s)$ with edges $E_s = \{\{i, j\} : j \in s_i\}$, and the cost to an agent i given s is defined

---

[4]We will sometimes use the shorthand $\ell_s(\cdot, \cdot)$ in place of $\ell_{G_s}(\cdot, \cdot)$.

[5]A *complete* graph/network refers to a network in which all pairs of nodes are directly linked to each other.

[6] A *star* refers to a minimally connected network in which there is one central node that is directly linked to the remaining $n - 1$ nodes. In a *center-sponsored star*, the cost of all $n - 1$ edges is borne by the center node. In a *peripheral-sponsored star*, each of the $n - 1$ peripheral agents bares the cost of building the edge connecting them to the center node.

to be

$$c_i(s) = d_i(s) = \sum_{j \in N} \ell_{G_s}(i, j). \tag{2.4}$$

The social cost for the UBBC game is defined as it is in the UC game; i.e., the sum of the costs incurred by the individual agents,

$$C(s) = \sum_{i \in N} c_i(s).$$

The remainder of this section includes some simple facts regarding Nash equilibria and efficient outcomes of the UBBC game. The first lemma is from Eshani et al. and identifies a sufficient condition for a joint strategy to be a Nash equilibrium.

Lemma 2.2. [13] *A UBBC strategy profile* $s = (s_1, \ldots, s_n)$ *that induces a network* G_s *without parallel edges and a diameter of at most* 2 *is a Nash equilibrium.*

The next two results identify properties of efficient outcomes for uniform instances of the UBBC game. Lemma 2.3 discerns the social cost of efficient networks while Proposition 2.4 shows that, for sufficiently sparse instances, efficient outcomes are necessarily networks with a diameter of two. The proofs of Lemma 2.3 and Proposition 2.4 can be found in the Appendix.

Lemma 2.3. *The social cost of any efficient strategy profile for a uniform UBBC instance with edge budgets* $k \geq 1$ *is*

$$2n(n - 1) - 2nk.$$

Proposition 2.4. *Every uniform UBBC instance with* $k < (n - 1)/2$ *has an efficient outcome with a diameter of* 2.

3. NIR UPPER BOUNDS

This section presents our results on upper bounding the NIR for the two network formation games. NIR upper bounds are established for the UC game in Section 3.1 (see Theorem 3.1) and for the UBBC game in Section 3.2 (see Theorem 3.7).

3.1. The Undirected Connections Game

We establish upper bounds on the NIR for the UC game for two regimes of α: when $\alpha < 1$ and $\alpha \geq 1$. These bounds are stated in Theorem 3.1, showing that inequality is independent of the number of agents when α is a constant. Our upper bound implies that when $\alpha > 3$, inequality scales linearly with α.

Theorem 3.1. *Upper bounds on the NIR for the UC game:*

1. When $\alpha < 1$*, the NIR is at most* $1 + \alpha < 2$.

2. *When $1 \leq \alpha < \infty$ is a constant (independent of n), the NIR is at most $\max\{2, (1+\alpha)/2\}$ in the limit as $n \to \infty$.*

We will prove the two parts of Theorem 3.1 separately in Lemma 3.2 (next) and Lemma 3.6 (following).

Lemma 3.2. *When $\alpha < 1$, the NIR for the UC games is at most $1 + \alpha < 2$.*

Proof. From Proposition 2.1 we know that whenever $\alpha < 1$, then the only Nash equilibrium is a complete graph, and every agent is adjacent to $n - 1$ other agents. In the complete graph, the usage costs incurred by every agent is the same, and the only disparity that can arise is due to the agents' construction costs. In the most extreme case, the min-cost agent does not buy any links and the max-cost agent buys $m > 0$ links. The inequality ratio is therefore

$$\frac{(n-1) + \alpha m}{n - 1} = 1 + \frac{\alpha m}{n - 1} \leq 1 + \alpha,$$

where m can be at most $n - 1$ (i.e., $m = n - 1$ when the max-cost agent buys all of its incident links). It is easy to see that the max-cost agent is not inclined to discard any of its edges because doing so would increase the usage cost by 1 but of yield a savings only $\alpha < 1$ in the construction cost, resulting in a net cost increase. □

Before addressing the second case of Theorem 3.1, we first need a couple of lemmas. In Lemma 3.3, we give an expression that captures the inequality ratio for an arbitrary Nash equilibrium of the UC game, and then show in Corollary 3.4 that this quantity is maximized in star networks. Then, in Lemma 3.5, we bound the inequality ratio of star network topologies.

Lemma 3.3 and Corollary 3.4 require additional notation. We refer to min-cost and max-cost agents as $\mu \in \text{argmin}_{i \in N} c_i(s)$ and $\chi \in \text{argmax}_{i \in N} c_i(s)$, respectively. Note that, in general, there may exist multiple min- and max-cost agents in any given Nash equilibrium, but in the following, we assume that μ and χ refer to the same min- and max-cost agents throughout the argument. For an agent i and a joint strategy s, let $T_s(i)$ denote a *shortest-path tree* rooted at i that is built from a breadth-first traversal of G_s beginning at node/agent i.[7] Edges of G_s that appear in $T_s(i)$ are referred to as *tree edges*, and those that do not are *nontree edges*. By construction, for any agent j that appears in layer $k \geq 0$ of $T_s(i)$, it follows that $\ell_s(i, j) = k$. Also, for any agent j appearing in layer k and every nontree edge $\{j, j'\}$, it follows (from the construction of $T_s(i)$) that agent j' must be in layer $k' \in \{k - 1, k, k + 1\}$.

For a pair of agents i, j, let $T_s(j; i)$ denote the subtree of $T_s(i)$ that is rooted at j such that, without loss of generality, all agents j' that are adjacent to j in G_s and are in layer $\ell_s(i, j) + 1$ of $T_s(i)$, are present in the first layer of the subtree $T_s(j; i)$. With a slight abuse of notation, we will use $T_s(j; i)$ to refer to the set of agents in the subtree as well as the subtree itself. Finally, let $\bar{T}_s(j; i)$ refer to the subtree (and its constituent agents) of $T_s(i)$

---

[7] Arguments using such trees are commonly used in the analysis of the PoA for network formation games; see, for example, Albers et al. [1].
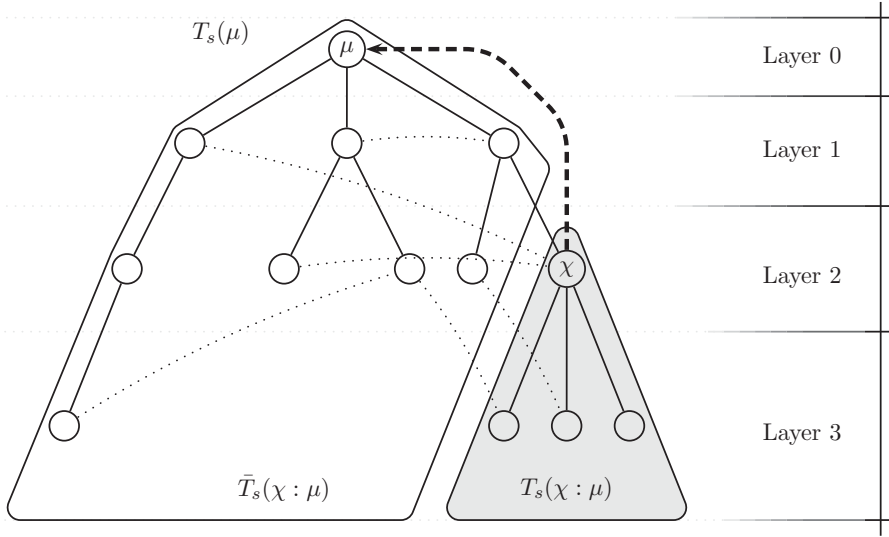
Figure 1 An induced shortest-path tree, $T_s(\mu)$, of a graph G_s rooted at a min-cost agent μ. Solid edges represent tree edges, and dotted edges represent nontree edges. Given a max-cost agent $\chi \neq \mu$ in layer $k = \ell_s(\mu, \chi)$, the tree $T_s(\chi; \mu)$ is the subtree rooted at χ with the stipulation that every i in layer $k + 1$ that is adjacent to χ in G_s is in $T_s(\chi; \mu)$. $\bar{T}_s(\chi; \mu)$ represents the part of $T_s(\mu)$ made up of agents/nodes that are not a part of $T_s(\chi; \mu)$.

that are not a part of $T_i(j; i)$. See Figure 1 for an illustration with $i = \mu$ and $j = \chi$ and solid (dotted) lines representing tree (nontree) edges.

Lemma 3.3. *Let μ and χ be min- and max-cost agents in a Nash equilibrium s. The inequality ratio is at most*

$$\text{IR}(s) = \frac{c_\chi(s)}{c_\mu(s)} \leq \frac{\alpha \left(1 + |s_\chi \cap T_s(\chi; \mu)|\right) + |\bar{T}_s(\chi; \mu)| - \ell_s(\chi, \mu) \cdot |T_s(\chi; \mu)| + d_\mu(s)}{d_\mu(s)},$$

(3.1)

where $d_\mu(s) = \sum_i \ell_s(\mu, i)$ is the distance-cost for agent μ.

Intuitively, the bound in Equation 3.1 follows from deriving upper- and lower-bounds on the costs of the max- and min-cost agents, respectively. The upper bound on the cost of the max-cost agent, χ, follows from considering a strategic deviation whereby χ builds a direct link to a min-cost agent μ and then drops all of its purchased links except those to agents $j \in s_\chi$ that are further from μ than χ. If the original strategy is a Nash equilibrium, then we know that χ's cost must be no higher than its cost given this new strategy.[8] The cost lower bound for a min-cost agent, μ, follows from simply assuming that μ did not purchase any edges (i.e., $c_\mu(s) \geq d_\mu(s)$).

Proof. Let χ and μ be a max- and a min-cost agent, respectively, in a Nash equilibrium strategy s, and let $T_s(\mu)$, $T_s(\chi; \mu)$, and $\bar{T}_s(\chi; \mu)$ be as defined above. Consider the strategy

---

[8]Our derivation of this upper bound on χ's cost extends a similar argument used by [1]. In particular, compare our Equation (3.2) below with Equation (2) in [1].

s'_χ for agent χ that is obtained from s_χ by dropping all of χ's edges to agents outside of $T_s(\chi; \mu)$ and buying an edge to μ (see Figure 1). Agent χ's cost under this new strategy $s' = (s'_\chi, s_{-\chi})$ is

$$c_\chi(s') \leq \underbrace{\alpha \left(1 + |s_\chi \cap T_s(\chi; \mu)|\right)}_{b_\chi(s')} + \underbrace{|\bar{T}_s(\chi; \mu)| - \ell_s(\mu, \chi) \cdot |T_s(\chi; \mu)| + d_\mu(s)}_{\text{upper-bound on } d_\chi(s')}. \quad (3.2)$$

Because s is a Nash equilibrium, we know that $c_\chi(s) \leq c_\chi(s')$, and the lemma follows from (3.2) and the fact that $c_\mu(s) \geq d_\mu(s)$.

Agent χ's usage cost, $b_\chi(s')$, simply reflects the construction of s'_χ – building an edge to μ and dropping all edges to nodes in s_χ that are outside of $T_s(\chi; \mu)$. The remainder of the proof is dedicated to establishing the upper bound on $d_\chi(s')$. For a set $X \subset N$, let $d_i^X(s) = \sum_{j \in X} \ell_s(i, j)$ denote agent i's usage cost to agents in X given the joint strategy s. We will bound $d_\chi(s')$ in terms of $d_\mu(s) = d_\mu^{T_s(\chi; \mu)}(s) + d_\mu^{\bar{T}_s(\chi; \mu)}(s)$. Notice that

$$d_\mu^{T_s(\chi; \mu)}(s) = \sum_{i \in T_s(\chi; \mu)} (\ell_s(\mu, \chi) + \ell_s(\chi, i))$$

$$= \ell_s(\mu, \chi) \cdot |T_s(\chi; \mu)| + \sum_{i \in T_s(\chi; \mu)} \ell_s(\chi, i)$$

$$= \ell_s(\mu, \chi) \cdot |T_s(\chi; \mu)| + d_\chi^{T_s(\chi; \mu)}(s),$$

since χ is on a shortest-path between μ and every agent $i \in T_s(\chi; \mu)$. Rearranging, we get

$$d_\chi^{T_s(\chi; \mu)}(s) = d_\mu^{T_s(\chi; \mu)}(s) - \ell_s(\mu, \chi) \cdot |T_s(\chi; \mu)|.$$

In s', χ's distance to agents in $T_s(\chi; \mu)$ is unchanged from s, so we can express χ's usage cost to agents in $T_s(\chi; \mu)$ under s' in relation to μ's distance to agents in $T_s(\chi; \mu)$ under s as

$$d_\chi^{T_s(\chi; \mu)}(s') = d_\mu^{T_s(\chi; \mu)}(s) - \ell_s(\mu, \chi) \cdot |T_s(\chi; \mu)|. \quad (3.3)$$

Next, we can derive $d_\chi^{\bar{T}_s(\chi; \mu)}(s')$ in terms of $d_\mu^{\bar{T}_s(\chi; \mu)}(s)$ in a very straightforward way. By including a link to μ in s'_χ and dropping all links to agents outside of $T_s(\chi; \mu)$, χ's usage cost to agents in $\bar{T}_s(\chi; \mu)$ is at most $|\bar{T}_s(\chi; \mu)|$ more than μ's usage cost to agents in $\bar{T}_s(\chi; \mu)$ under s; i.e.,

$$d_\chi^{\bar{T}_s(\chi; \mu)}(s') \leq |\bar{T}_s(\chi; \mu)| + d_\mu^{\bar{T}_s(\chi; \mu)}(s). \quad (3.4)$$

Putting (3.3) and (3.4) together, we get

$$d_\chi(s') = d_\chi^{T_s(\chi; \mu)}(s') + d_\chi^{\bar{T}_s(\chi; \mu)}(s')$$

$$\leq d_\mu^{T_s(\chi; \mu)}(s) - \ell_s(\mu, \chi) \cdot |T_s(\chi; \mu)| + |\bar{T}_s(\chi; \mu)| + d_\mu^{\bar{T}_s(\chi; \mu)}(s)$$

$$= d_\mu(s) - \ell_s(\mu, \chi) \cdot |T_s(\chi; \mu)| + |\bar{T}_s(\chi; \mu)|,$$

where we substituted $d_\mu(s) = d_\mu^{T_s(\chi; \mu)}(s) + d_\mu^{\bar{T}_s(\chi; \mu)}(s)$ in the last line. This completes the proof. $\qquad \square$

Corollary 3.4. *Among the Nash equilibrium strategies that arise when the edge construc-tion cost is a constant $1 \leq \alpha < \infty$, the star topology maximizes the inequality ratio.*

Proof. We will show that the inequality ratio established in Lemma 3.3 is maximized in a star network topology. Recall Equation (3.1), bounding the inequality ratio of an arbitrary Nash equilibrium strategy s for the UC game, repeated here for convenience:

$$\text{IR}(s) \leq \frac{\alpha(1 + |s_\chi \cap T_s(\chi;\mu)|) + |\bar{T}_s(\chi;\mu)| - \ell_s(\chi,\mu) \cdot |T_s(\chi;\mu)| + d_\mu(s)}{d_\mu(s)}. \quad (3.1)$$

Let $x = |T_s(\chi;\mu)|$ (ergo, $|\bar{T}_s(\chi;\mu)| = n-x$) and $x_0 = 1 + |s_\chi \cap T_s(\chi;\mu)| \leq x$. Substituting

$$d_\mu(s) = d_\mu^{T_s(\chi;\mu)}(s) + d_\mu^{\bar{T}_s(\chi;\mu)}(s)$$

$$= \ell_s(\chi,\mu)x + d_\chi^{T_s(\chi;\mu)}(s) + d_\mu^{\bar{T}_s(\chi;\mu)}(s),$$

we can rewrite (3.1) as

$$\text{IR}(s) \leq \frac{\alpha x_0 + (n-x) - \ell_s(\chi,\mu)x + d_\mu(s)}{d_\mu(s)}$$

$$= \frac{\alpha x_0 + n - x - \ell_s(\chi,\mu)x + d_\mu^{T_s(\chi;\mu)}(s) + d_\mu^{\bar{T}_s(\chi;\mu)}(s)}{d_\mu^{T_s(\chi;\mu)}(s) + d_\mu^{\bar{T}_s(\chi;\mu)}(s)}$$

$$= \frac{\alpha x_0 + n - x - \ell_s(\chi,\mu)x + \ell_s(\chi,\mu)x + d_\chi^{T_s(\chi;\mu)}(s) + d_\mu^{\bar{T}_s(\chi;\mu)}(s)}{\ell_s(\chi,\mu)x + d_\chi^{T_s(\chi;\mu)}(s) + d_\mu^{\bar{T}_s(\chi;\mu)}(s)}$$

$$= \frac{\alpha x_0 - n - x + d_\chi^{T_s(\chi;\mu)}(s) + d_\mu^{\bar{T}_s(\chi;\mu)}(s)}{\ell_s(\chi,\mu)x + d_\chi^{T_s(\chi;\mu)}(s) + d_\mu^{\bar{T}_s(\chi;\mu)}(s)}. \quad (3.5)$$

Notice that in Equation (3.5), the distance between μ and χ in G_s appears only in the denominator. Therefore, toward our aim of maximizing (3.5), we can infer that $\ell_s(\chi,\mu) = 1$. Hence, μ and χ are necessarily adjacent in a Nash equilibrium s that maximizes the inequality ratio. With μ and χ adjacent, we can substitute

$$d_\mu^{\bar{T}_s(\chi;\mu)}(s) = d_\chi^{\bar{T}_s(\chi;\mu)}(s) - |\bar{T}_s(\chi;\mu)| = d_\chi^{\bar{T}_s(\chi;\mu)}(s) - n + x$$

and

$$d_\chi^{T_s(\chi;\mu)}(s) = d_\mu^{T_s(\chi;\mu)}(s) - |T_s(\chi;\mu)| = d_\mu^{T_s(\chi;\mu)}(s) - x,$$

giving us

$$\text{IR}(s) \leq \frac{\alpha x_0 + d_\chi(s) - 2x - n}{d_\mu(s) + 2x - n}. \quad (3.6)$$

This leaves us with two approaches to identify topologies that maximize the inequality ratio: we can maximize x_0 (i.e., by setting $x_0 = n-1$), or we can minimize x (i.e., set $x = 1$). Both approaches imply star network topologies (the center- and peripheral-sponsored stars, respectively), completing the proof. $\qquad\square$

The next lemma establishes asymptotically tight upper bound on the inequality ratio for star network topologies when the edge cost $1 \leq \alpha < \infty$.

Lemma 3.5. *As the number of agents grows toward infinity, the maximal inequality in a star topology for the UC network formation game with a constant edge cost $1 \leq \alpha < \infty$ is* $\max\{2, (1 + \alpha)/2\}$.

The proof of Lemma 3.5 is straightforward, so we provide only a sketch here. The full proof can be found in the Appendix.

Proof Sketch. We consider two star topologies: the center-sponsored star strategy and the peripheral-sponsored star strategy. It is easy to see that the inequality in these two star strategies dominate that of all other star strategies (see the full proof in the Appendix for details).

In the center-sponsored star strategy, s^c, the max-cost agent χ is in the center and the min-cost agent μ is in the periphery. Therefore, the inequality ratio is

$$\text{IR}(s^c) = \frac{c_\chi(s^c)}{c_\mu(s^c)} = \frac{(\alpha + 1)(n - 1)}{2n - 3}.$$

In the peripheral-sponsored star strategy, s^p, the min-cost agent μ is in the center and the max-cost agent χ is in the periphery. Hence, the inequality ratio is

$$\text{IR}(s^p) = \frac{c_\chi(s^p)}{c_\mu(s^p)} = \frac{\alpha + 2n - 3}{n - 1}.$$

As $n \to \infty$, $\text{IR}(s^c)$ approaches $(1+\alpha)/2$ and $\text{IR}(s^p)$ approaches 2. Therefore, the inequality ratio for a star network is $\max\{2, (1 + \alpha)/2\}$ as n grows to infinity. $\square$

Finally, we are ready to prove the second case of Theorem 3.1, which will follow directly from Lemmas 3.3 and 3.5 and Corollary 3.4.

Lemma 3.6. *When $1 \leq \alpha < \infty$ is a constant (independent of n), the NIR for the UC game is at most* $\max\{2, (1 + \alpha)/2\}$ *in the limit as $n \to \infty$.*

Proof. From Corollary 3.4 we know that, among Nash equilibrium strategies, the star maximizes the inequality ratios and Lemma 3.5 provides the desired IR upper bound for stars. $\square$

3.2. The Undirected Bounded Budget Connections Game

For the UBBC game, the NIR is at most two (Theorem 3.7). This bound holds for uniform instances of the game, in which every agent has the same edge budget, as well as for general cases in which agents can have heterogenous budgets. This upper bound is tight, demonstrating that inequality at equilibrium in the UBBC game is independent of *ex ante* inequalities in edge endowments.

Theorem 3.7. *The NIR for the UBBC game is at most* 2.

The proof of Theorem 3.7 (following) proceeds as follows: we argue that in any strategy profile where the agents with the maximum and minimum costs are adjacent, the inequality ratio must be strictly less than two. If, the min-cost and max-cost agents are *not* adjacent in a Nash equilibrium, it must be the case that neither of them stand to reduce their individual costs by switching to a strategy that includes a link with the other agent, because doing so would mean that the strategy is not a Nash equilibrium in the first place. This implies that either

1. The inequality ratio is already less than 2; or
2. A maximum cost agent switching to a strategy that includes a link to a minimum cost agent would lead them to a *higher* cost (i.e., by disconnecting the network).

However, by Lemma 3.8, case 2 cannot be a Nash equilibrium, so the NIR can be at most 2.

Lemma 3.8. *Suppose $s = (s_1, \ldots, s_n)$ is a UBBC strategy in which there exist (distinct) agents $x, y, z \in N$ with $x \in s_z$, $y \notin s_z$, and $z \notin s_y$ such that*

1. *Every $x \to y$ path in G_s contains agent z as an intermediate node, and*
2. *If z were to swap out x for a connection with agent y, then the network will become disconnected (i.e., the edge $\{z, x\}$ is a bridge).*

Then s cannot be a Nash equilibrium.

The proof of Lemma 3.8 employs the main result of [26] with regard to the structure of networks in asymmetric swap-equilibrium. The asymmetric swap-equilibrium holds for a strategy s when, for every agent $i \in N$ and every deviation s_i' that differs from s_i with the addition, removal, or swap of a single edge, we have $c_i(s_i, s_{-i}) \leq c_i(s_i', s_{-i})$. The asymmetric swap-equilibrium is a weaker equilibrium concept than the Nash equilibrium for the UBBC, so their result (restated in Theorem 3.9 below) carries over to Nash equilibrium for the UBBC. Recall that a graph is k-edge-connected if the removal of any $k - 1$ edge does not disconnect it. A k-edge-connected component of a graph G is a maximal subgraph $G' \subset G$ that is k-edge-connected.

Theorem 3.9. [26] *Every network in an asymmetric swap-equilibrium has at most one 2-edge-connected component.*

Put another way, Theorem 3.9 states that a Nash equilibrium strategy s will not induce a graph G_s in which the removal of a single edge will split the graph into two components that are both 2-edge-connected; i.e., the two components cannot both contain cycles.

Proof of Lemma 3.8. Let s be a strategy with agents $x, y, z \in N$ that satisfy the given conditions and assume (toward a contradiction) that s is a Nash equilibrium. Denote by X the component of the network containing x that would result from the removal of x from agent z's strategy, and let $Y = N \setminus X \cup \{z\}$ be the set of the remaining agents (not in X) also excluding agent z (see Figure 2). Recall that every agent i contributes at least one edge to G_s (since $k_i \geq 1$) so the component X must contain a cycle. Hence, by Theorem 3.9, we know that the subgraph induced by agents in Y must not contain a cycle.
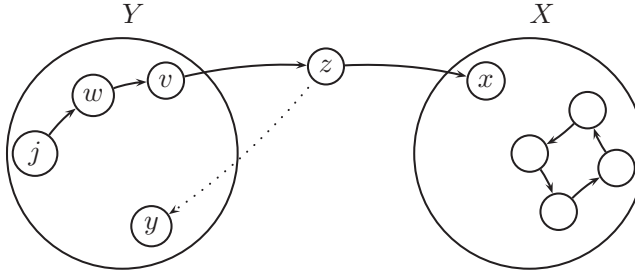
Figure 2 A schematic of the conditions expressed in Lemma 3.8.

Let G_Y denote the subgraph induced by the agents in Y. If $\exists j \in Y$ such that $k_j > 1$ then either $|s_j| < k_j$ and j is not playing a best response (because they could build an edge to an nonneighboring agent and decrease their cost) or G_Y will contain a cycle; both scenarios contradicting our assumption that s is a Nash equilibrium. Therefore, assume that the subgraph G_Y is a tree, and that $k_j = 1$ for all $j \in Y$. Clearly, for every path $p = (y_1, y_2, \ldots, y_m = z)$ that begins at an agent $y_1 \in Y$ and ends at z, we have $y_{i+1} \in s_{y_i}$ (to be otherwise would imply that $k_j > 1$ for some $j \in Y$).

Let $j \in Y$ be a leaf in G_Y that is furthest from z. Let $w \in s_j$ be j's parent, and let $v \in s_w$ be j's grandparent. If $w = z$, then we are done because this contradicts our assumption that y and z are not adjacent in G_s. So assume that $w \neq z$. Let D_w be w's descendants in G_Y (i.e., the set of agents for whom w lies on their unique path to z), and let $C_w \subseteq D_w$ be w's children (i.e., the set of agents in D_w that are adjacent to w). Consider a deviation by j to the strategy $s'_j = \{v\}$ that replaces j's link to their parent with a link to their grandparent. This deviation will increase j's distance to w and all nodes in $C_w \setminus \{j\}$ by one and decrease j's distance to every node in $N \setminus C_w \setminus \{j\}$ by one. If $|C_w| < |N \setminus C_w| - 1$ then agent j will benefit from switching to strategy s'_j, contradicting our assumption that s is a Nash equilibrium.

Suppose that $|C_w| \geq |N \setminus C_w| - 1$, which implies that $|C_w| \geq (n-1)/2$. Let $u \in s_v$ be w's grandparent in G_Y. If there is another agent w' with $v \in s_{w'}$, then this agent can swap their link to v with a link to w and receive a strictly lower cost, which implies that s was not a Nash equilibrium. If there is no such w', then w can swap its link to v with a link to its grandparent u and receive a strictly lower cost. This again implies that s was not a Nash equilibrium. This completes the proof.

We are now ready to attend to the proof of Theorem 3.7. We will use the shorthand $G - \{i, j\}$ to denote the graph G with the edge $\{i, j\}$ removed. $\qquad\square$

Proof of Theorem 3.7. Consider (toward a contradiction) a strategy profile s that is a Nash equilibrium in which $c_{\max}(s) \geq 2 \cdot c_{\min}(s)$; i.e., a strategy s such that the inequality ratio is at least 2. Notice that in this strategy an agent χ with $c_\chi(s) = c_{\max}(s)$ cannot be directly connected to an agent μ with $c_\mu(s) = c_{\min}(s)$, because if it were, then χ would be connected to the $n - 2$ other agents via μ's shortest paths for a cost that is at most $n - 2$ more than the cost μ is subjected to. That is, if $\chi \in s_\mu$ or $\mu \in s_\chi$ then

$$c_\chi(s) \leq (n - 2) + c_\mu(s) < 2 \cdot c_\mu(s). \tag{3.7}$$

The strict inequality in (3.7) comes from the fact that $c_\mu(s) > n - 2$ since it is impossible for μ to connect to $n - 1$ other agents for a cost any less than $n - 1$. Therefore, in order for $c_\chi(s) \geq 2 \cdot c_\mu(s)$, it must be the case that $\mu \notin s_\chi$ and $\chi \notin s_\mu$. This implies that $c_\chi(s)$ must be no greater than $(n - 2) + c_\mu(s)$, because switching to a strategy s'_χ that includes a link to agent μ would ensure as much, provided that switching to s'_χ does not disconnect the network.

Suppose that $c_\chi(s) > (n - 2) + c_\mu(s)$ and if χ were to switch to a strategy s'_χ that is obtained by swapping out some $x \in s_\chi$ for a link to μ results in the network $G_{(s'_\chi, s_{-\chi})}$ becoming disconnected. Then it must be that every $x \to \mu$ path in G_s includes agent χ, in which case it follows that the component of $G_s - \{\chi, x\}$ including x contains a cycle because every agent's budget affords them at least one edge. However, by Lemma 3.8, such a strategy s cannot be a Nash equilibrium, contradicting our assumption that s is a Nash equilibrium. This completes the proof.   □

The upper bound established by Theorem 3.7 is asymptotically tight. This can be observed in a "star-like" network in which every agent has a link to the min-cost agent μ, and the degree of the max-cost agent χ is bounded by a constant independent of the number of agents, n.[9] In the asymptotic limit, as the number of agents tends toward infinity, the inequality ratio between χ and μ equals 2.

4. EQUALITY, EQUILIBRIUM, AND EFFICIENCY

With the upper bounds on inequality in Nash equilibrium strategies established, we are now in a position to address the relationship between (in)equality, equilibrium, and efficiency. Although our focus here is on a specific metric of inequality (the inequality ratio) and how it relates to efficiency in a particular setting (the network formation games), the connection between inequality and efficiency more generally is a matter of considerable interest among economic analysts, researchers, and policy makers (cf. [6, 7, 12, 8, 9]), as well as the general public [31, 29]. As in Section 3, the presentation of this section is divided into two parts, with the UC game analyzed in 4.1 (see Theorem 4.1) followed by the UBBC in Section 4.2 (see Theorem 4.5).

4.1. The Undirected Connections Game

Theorem 4.1 summarizes our results on the relationship between equality, efficiency, and Nash equilibrium for the UC game for the three regimes of α identified in Proposition 2.1.

Theorem 4.1. *The relationship between inequality, Nash equilibrium, and efficiency in the UC game:*

1. *When $\alpha < 1$, there exist efficient Nash equilibrium strategies that maximize the inequality ratio and other efficient Nash equilibrium strategies that achieve cost equality in the limit as $n \to \infty$.*

---

[9]For a concrete example, consider a uniform instance of the UBBC game in which $k_i = 1$ for all $i \in N$, and the Nash equilibrium strategy that forms a star network with one "extra" edge linking a pair of peripheral agents.

2. *When $1 \leq \alpha < 2$, no Nash equilibrium strategy is also efficient. However, among Nash equilibrium strategies, there exist some that maximize the inequality ratio, whereas others are egalitarian in the limit as $n \to \infty$.*
3. *When $2 \leq \alpha < \infty$, there exist Nash equilibrium strategies that are*

   - *both efficient and egalitarian in the limit as $n \to \infty$; and*
   - *efficient with maximal inequality in the limit as $n \to \infty$.*

We will prove the three parts of Theorem 4.1 individually in the next three lemmas.

Lemma 4.2. *When $\alpha < 1$, there exist efficient Nash equilibrium strategies for the UC game that maximize the inequality ratio, and other efficient Nash equilibrium strategies that achieve cost equality in the limit as $n \to \infty$.*

Proof. By Proposition 2.1, when $\alpha < 1$, both Nash equilibrium strategies and socially efficient strategies result in the formation of the complete network. From this fact and Theorem 3.1, it follows that there exist socially efficient strategies that maximize the inequality ratio. Therefore, we only need to show that there exist equilibrium strategies that achieve cost equality.

Equality among all agents would follow if every agent builds exactly $(n-1)/2$ edges, which can only occur when n is odd. When, n is even, then $(n-1)/2$ is nonintegral; so the closest we can get to equality calls for a strategy s in which the max-cost agents each build $\lceil (n-1)/2 \rceil$ edges while the min-cost agents get away with building one fewer edge each. In this case, a min-cost agent's cost is $c_{\min}(s) = c_{\max}(s) - \alpha$, and the inequality ratio is

$$
\begin{aligned}
\mathrm{IR}(s) &= \frac{c_{\max}(s)}{c_{\max}(s) - \alpha} \\
&= \frac{(n-1) - \lceil (n-1)/2 \rceil \alpha}{(n-1) - (\lceil (n-1)/2 \rceil - 1)\alpha} \\
&= 1 - \frac{\alpha}{\alpha + \alpha \lfloor (1-n)/2 \rfloor + n - 1}.
\end{aligned}
\tag{4.1}
$$

As $n \to \infty$, Equation (4.1) approaches 1. Hence, inequality vanishes as n increases, concluding the proof. $\square$

Turning to the second item of Theorem 4.1, which addresses the case that $\alpha \in [1, 2)$, we first note that Proposition 2.1 already implies that no Nash equilibrium strategy is also an efficient strategy. This is because the Nash equilibrium strategies for this regime of α correspond to networks with star topologies while efficient strategies correspond to complete networks in which every possible edge is present. Furthermore, by Corollary 3.4 we know that the upper bound on the inequality ratio for this regime of α is achieved in a star topology. Hence, to prove the second item of Theorem 4.1 we need only address the existence of equilibrium strategies for which there is cost equality among the agents.

Lemma 4.3. *When $1 \leq \alpha < 2$, the UC game admits a Nash equilibrium strategy that is egalitarian in the limit as $n \to \infty$.*

Proof. Consider a joint strategy s inducing a star network topology among the $|N| = n$ agents in which the central agent $c \in N$ buys $k = 1 + \lfloor \frac{n-2}{\alpha} \rfloor$ edges $\{c, i\}$ to agents $N_{\neg b} \subset N$; all the remaining edges $\{j, c\}$ are paid for by their respective peripheral agent $j \in N_b = N \setminus N_{\neg b} \cup \{c\}$. (The sets N_b and $N_{\neg b}$ partition N so that all agents that buy edges appear in N_b and all those who do not buy any edges appear in $N_{\neg b}$.) In this strategy, the agents $i \in N_{\neg b}$ all incur a cost $c_i(s) = 2n - 3$, all agents $j \in N_b$ incur a cost of $c_j(s) = 2n - 3 + \alpha$, and c incurs a cost $c_c(s) = n - 1 + \alpha k \leq 2n - 3 + \alpha$. Since $c_i(s) < c_c(s) \leq c_j(s)$, the inequality ratio is:

$$\mathrm{IR}(s) = \frac{c_j(s)}{c_i(s)} = \frac{2n - 3 + \alpha}{2n - 3} = 1 + \frac{\alpha}{2n - 3}, \tag{4.2}$$

which approaches 1 in the limit as $n \to \infty$. It is easy to see that s is in fact a Nash equilibrium because no agent stands to reduce its cost by deleting an edge (because doing so would disconnect the network, prompting an infinite cost) and, at the cost of $\alpha \geq 1$ per edge, the addition of any edges will decrease the agent's usage cost by only 1 (per edge), yielding no net decrease in the agent's total cost.   □

Finally, the third item of Theorem 4.1 addresses the regime when α is a constant greater than or equal to 2. In this regime, we show that there exist some efficient Nash equilibrium strategies that are egalitarian and other efficient Nash equilibrium strategies that achieve maximal inequality.

Lemma 4.4. *When $2 \leq \alpha < \infty$, there exist Nash equilibrium strategies in the UC game that are both efficient and egalitarian in the limit as $n \to \infty$; and there are Nash equilibrium strategies that are both efficient and achieve maximal inequality.*

Proof. Recall from Proposition 2.1 that, when $\alpha \geq 2$, the star is an efficient Nash equilibrium strategy. Thus, the lemma will follow if we can show that some star strategies can support maximal inequality while others can support equality.

To show the first part of the lemma (the existence of an egalitarian strategy that yields to a star network topology), we can invoke the strategy constructed in the proof of Lemma 4.3. Recall that this strategy calls for the central agent $c \in N$ to buy $1 + \lfloor \frac{n-2}{\alpha} \rfloor$ edges to peripheral agents $N_{\neg b} \subset N$, and all remaining agents $j \in N \setminus N_{\neg b} \cup \{c\}$ buy a single edge to c. The inequality ratio for this strategy is expressed by Equation 4.2, which approaches 1 in the limit as $n \to \infty$, and is, hence, egalitarian.

The second part of the lemma (the existence of a star-topology-yielding strategy that achieves the inequality upper bound established in Theorem 3.1) follows immediately from Corollary 3.4, which states that inequality in the UC game is maximized in networks with a star topology.   □

4.2. The Undirected Bounded Budget Connections Game

This section establishes the relationship between efficient Nash equilibrium strategies and (in)equality in the UBBC game with uniform budgets. We find that when the budget k is sufficiently small ($k = O(1)$), we can construct an efficient Nash equilibrium strategy in which inequality is maximized; and when the edge budget is sufficiently large ($k \geq 2\sqrt{n-1}$), then we show the existence of efficient Nash equilibrium strategies in which
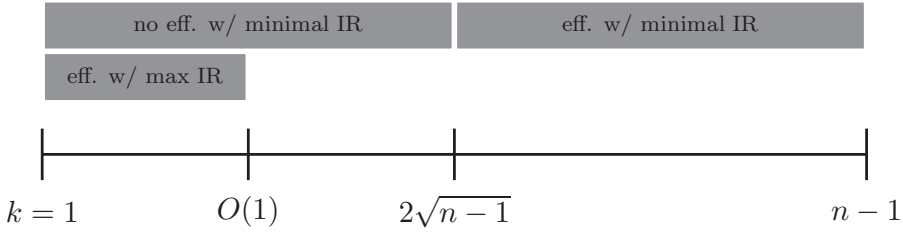
Figure 3 A summary of Theorem 4.5, characterizing the existence of Nash equilibrium strategies in the UBBC game with respect to efficiency and the inequality ratio for different uniform budgets, k.

there is cost equality among all of the agents. Finally, we prove the nonexistence of efficient Nash equilibrium strategies with egalitarian costs when $k < 2\sqrt{n-1}$. These findings are formally presented in Theorem 4.5 and summarized in Figure 3.

Theorem 4.5. *The relationship between inequality, Nash equilibrium, and efficiency in the UBBC game with uniform budgets k:*

1. *When $k \geq 1$ is a constant (independent of n), there exists an efficient Nash equilibrium strategy that achieves the NIR upper bound of 2 established in Theorem 3.7.*
2. *When $k \geq 2\sqrt{n-1}$, there exist efficient Nash equilibrium strategies that are egalitarian.*
3. *When $1 \leq k < 2\sqrt{n-1}$, there does not exist an efficient Nash equilibrium strategy that is egalitarian.*

Lemma 4.6. *In the UBBC game with constant (independent from n) uniform edge budgets $k \geq 1$, there exists an efficient Nash equilibrium strategy that achieves the NIR upper bound of 2.*

Proof. Fix k and let $N = \{0, 1, \ldots, n-1\}$ denote the set of strategic agents. We will construct an efficient Nash equilibrium strategy s as follows (see Figure 4):

1. $\forall i \in \{0, 1, \ldots, k\}$, set $s_i = \{i+1, i+2, \ldots, i+k\}$.
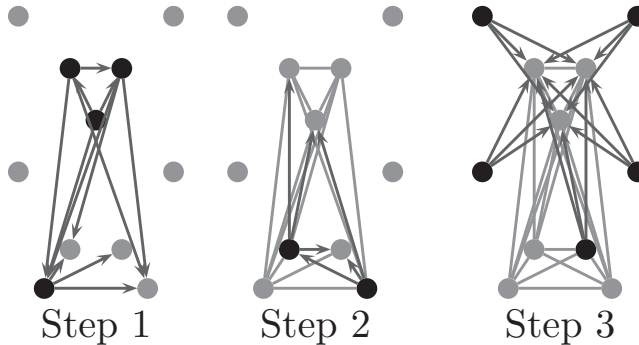
Figure 4 Example of the three parts to the strategy constructed in the proof of Lemma 4.6. Dark nodes and edges are those assigned in the specified step and arrows are meant to convey link "ownership."

2. $\forall i \in \{k + 1, k + 2, \ldots, 2k - 1\}$, set $s_i = \{0, 1, \ldots, i - k - 1, i + 1, i + 2, \ldots, 2k\}$.

3. $\forall i \in \{2k, 2k + 1, \ldots, n - 1\}$, set $s_i = \{0, 1, \ldots, k - 1\}$.

The strategy s produces a network that has a diameter of 2. Thus, per Lemma 2.2, it is a Nash equilibrium.

To show that s is socially efficient, we must determine its social cost. We can partition the agents N into three sets that correspond to three "tiers" of costs:

- Agents in $N_a = \{0, 1, \ldots, k - 1\}$ are each directly connected to every other agent, so $c_i(s) = n - 1$ for all $i \in N_a$. Agents in N_a incur the minimal cost among N.
- Agents in $N_b = \{2k + 1, 2k + 2, \ldots, n - 1\}$ are each directly linked to every agent in N_a and two hops away from every other agent, so $c_i(s) = 2(n - 1) - k$ for all $i \in N_b$. Agents in N_b incur the maximal cost among N.
- Agents $N_c = \{k, k + 1, \ldots, 2k\}$ each have a direct link to every agent in N_a and N_c and a two-hop distance to agents in N_b, so $c_i(s) = 2(n - k - 1)$ for all $i \in N_c$.

The cardinality of these partitions are $|N_a| = k$, $|N_b| = n - 2k - 1$, and $|N_c| = k + 1$, and the social cost is

$$C(s) = [k(n - 1)] + [(n - 2k - 1)(2(n - 1) - k)] + [(k + 1)(2(n - k - 1))]$$
$$= 2(n^2 - n - kn).$$

By Lemma 2.3, this is the minimal social cost in a uniform game with size-k budgets. Therefore, s is socially efficient.

The inequality ratio is between a min-cost agent $i \in N_a$ and a max-cost agent $j \in N_b$;

$$\text{IR}(s) = \frac{c_j(s)}{c_i(s)} = \frac{2(n - 1) - k}{n - 1} = 2 - \frac{k}{n - 1}. \tag{4.3}$$

Because (4.3) approaches 2 as $n \to \infty$, this strategy maximizes the inequality ratio established for the UBBC in Theorem 3.7. Hence, s is an efficient Nash equilibrium strategy that achieves maximal inequality. $\square$

The next lemma establishes the second part of Theorem 4.5, showing that when edge budgets are sufficiently large, we can find an efficient Nash equilibrium strategy that achieves perfect equality among the agents. But first we need the following fact (see the Appendix for a proof):

Fact 4.7. *A δ-regular graph of diameter two can have at most $n = \delta^2 + 1$ nodes.*

Lemma 4.8. *In the UBBC game with uniform edge budgets $k \geq 2\sqrt{n - 1}$, there exist efficient Nash equilibrium strategies that are egalitarian.*

Proof. For a strategy s to be efficient and egalitarian, it must be the case that G_s has a diameter of at most 2, and all agents have common one- and two-hop neighborhood sizes. When $k \geq (n - 1)/2$, then this is trivially satisfied by the complete graph. Satisfying both of these properties for $k < (n - 1)/2$ requires that G_s be $2k$-regular with a diameter of 2. From Fact 4.7, we know that such graphs are possible only if $k \geq 2\sqrt{n - 1}$. $\square$

Finally, we turn to the third part of Theorem 4.5, which establishes the nonexistence of efficient Nash equilibrium strategies that are egalitarian when edge budgets are below $2\sqrt{n-1}$. The nonexistence follows as a corollary of Fact 4.7 and Proposition 2.4 by observing that when $k < 2\sqrt{n-1}$, there are not enough edges to create a regular, diameter-2 network.

Corollary 4.9. *When* $1 \leq k < 2\sqrt{n-1}$*, there does not exist an efficient Nash equilibrium strategy that is egalitarian.*

5. DISCUSSION AND CONCLUSIONS

This paper examines inequality in two simple models of strategic network formation: the Undirected Connections (UC) game [14] and the Undirected Bounded Budget Connections (UBBC) game introduced by [13]. To this end, we introduce the Nash Inequality Ratio (NIR) as an instrument to quantify the level of inequality that can exist between a pair of agents in a Nash equilibrium. Upper bounds on the NIR are established for both games, and we show that these bounds are tight.

Our analyses of these games reveal an interesting relationship between scarcity (expressed by lower budgets in the UBBC model and high edge costs in the UC model) and inequality. In the UBBC game, we found that when edges are scarce, the inequality upper bound is attainable in Nash equilibrium; but when edges are more plentiful, the highest level of inequality is not sustainable in equilibrium. A similar correspondence is observed in the UC model, in which higher edge costs can support maximal inequality in equilibrium. With respect to efficiency, we find that the two games behave differently from one another: in the UC game, scarcity can support efficient Nash equilibrium strategies with either egalitarian costs or maximal inequality; but the UBBC game supports only efficient Nash equilibrium strategies that are not egalitarian when edges are scarce.

The NIR affords a wide-angle lens through which inequality can be analyzed, characterizing the *extent* to which the costs incurred by a pair of agents can differ in equilibrium. It does not, however, provide a clear view on the *distribution* of inequality among populations of agents, nor does it identify the inequality ratio of a "typical" Nash equilibrium strategy. Indeed, in our analysis of both network formation games we often relied on the inequality furnished by networks with a star topology in which there is a single, central agent that incurred the lowest (or sometimes highest) cost while the remaining $n-1$ agents all incurred the highest (or lowest) cost. In star networks, the inequality between a pair of agents selected uniformly at random is, with a high probability, nonexistent since neither agent will likely be the central agent. One interesting direction for future research that extends the analysis presented here would be to derive upper bounds on the expected inequality ratio between randomly selected pairs of agents. A related problem is to bound the inequality ratio of a randomly chosen Nash equilibrium strategy, or from the set of equilibrium strategies that result from a particular game dynamic such as best or approximate-best response.

We believe that an analysis of the NIR will provide interesting insight into many other games beyond the two network formation games analyzed here. As with the PoA, the NIR can be used to establish a bound on the "price" of strategic behavior by answering the questions: To what extent can the speed of costs/benefits be found in equilibrium outcomes of distributed decision making by self-interested agents? It may be anarchy, but is it fair?

ACKNOWLEDGMENTS

FUNDING

APPENDIX: OMITTED PROOFS

Lemma 2.3. *The social cost of any efficient strategy profile for a uniform UBBC instance with edge budgets $k \geq 1$ is*

$$2n(n-1) - 2nk. \tag{A.1}$$

Proof. Let $C_{\min}(n, k)$ be the social cost of an efficient outcome for a uniform UBBC instance with n agents, each with a budget of k edges. When $k = k_{\max} = (n-1)/2$, the complete graph is the efficient outcome with a social cost of $C_{\min}(n, k_{\max}) = n(n-1)$. Because an individual edge must lie on at least two shortest paths (e.g., the edge $\{i, j\}$ is on both the $i \rightarrow j$ and $j \rightarrow i$ paths), starting from the complete network, every edge removal increases the social cost by at least 2. Therefore, decreasing the edge budget by one reduces the social cost by at least $2n$. Hence, we can express $C_{\min}(n, k)$ by

$$C_{\min}(n, k) = n(n-1) + 2n(k_{\max} - k).$$

Substituting $k_{\max} = (n-1)/2$ gives

$$C_{\min}(n, k) = n(n-1) + 2n\left(\frac{n-1}{2} - k\right),$$

which can be rearranged into Equation (A.1), completing the proof. $\square$

Proposition 2.4. *Every uniform UBBC instance with $k < (n-1)/2$ has an efficient outcome with a diameter of 2.*

Proving Propositon 2.4 relies on the following lemma:

Lemma A.1. *The social costs of all diameter-2 topologies with a fixed number of $m \leq \frac{n(n-1)}{2}$ non-overlapping edges are equivalent.*

Proof. Suppose that s is a joint strategy profile that induces a diameter-2 network G_s with m distinct (i.e., nonparallel) edges. Because G_s has a diameter of 2, we can express the cost incurred by each agent $i \in N$ in terms of their degree and the cardinality of their one- and two-hop neighborhoods. Let $N_x(G_s, i)$ and $d(G_s, i)$ denote i's x-hop neighborhood and

degree, respectively, in G_s. The cost to i can be expressed as

$$c_i(s) = d(G_s, i) + 2|N_2(G_s, i) \setminus N_1(G_s, i)| = d(G_s i) + 2\left(|N_2(G_s, i)| - |N_1(G_s, i)|\right)$$
$$= d(G_s, i) + 2\left(|N_2(G_s, i)| - d(G_s, i) - 1\right) = 2|N_2(G_s, i)| - d(G_s i) - 2$$
$$= 2n - d(G_s, i) - 2 \tag{A.2}$$

The substitution of $|N_2(i)|$ for n in the last line is a requirement of our assumption that the network G_s has a diameter of 2.

With the cost to an individual agent in a diameter-2 network established in Equation (A.2), we can turn our attention to the social cost. Let $C(n, m)$ denote the social cost of a diameter-2 network G with n nodes and m edges.

$$C(n, m) = \sum_{i \in N} (2n - d(G, i) - 2) = 2n^2 - 2n - 2m. \tag{A.3}$$

The lemma follows from the fact that Equation (A.3) depends only on the diameter-2 assumption, the number of nodes, and the number of edges. $\qquad\square$

Proof of Proposition 2.4. From Lemma 2.3 we know that the social cost of an efficient outcome for a uniform UBBC instance is given by Equation (A.1), and from Lemma A.1 we know that the social cost of any diameter-2 network with m edges is given by Equation (A.3). A uniform UBBC instance with k edges per agent induces a network with $m = nk$ edges. Substituting for m and rearranging shows that both Equations (A.1) and (A.3) are equivalent. The requirement that $k < (n-1)/2$ is a consequence of the fact that when $k \geq (n-1)/2$, the efficient outcome is the complete graph, which has a diameter of 1. $\qquad\square$

Lemma 3.5. *As the number of agents grows toward infinity, the maximal inequality in a star topology for the UC network formation game with a constant edge cost $1 \leq \alpha < \infty$ is $\max\{2, (1 + \alpha)/2\}$.*

Note that the notation for the following proof of Lemma 3.5 differs from that used in the proof sketch given in Section 3.1. This is because the following proof must address more than just the center- and peripheral-sponsored stars as was done in the sketch.

Proof. Let s be a strategy profile that produces a star topology rooted at agent c. Suppose that $k \in \{0, 1, \ldots, n-1\}$ is the number of edges that c purchases; i.e., $|s_c| = k$. When $k = 0$, we get the *peripheral-sponsored star*, and when $k = n - 1$, we get the *center-sponsored star*. If $k < n - 1$, then $\exists j \in N$ such that $c_j(s) = 2n - 3 + \alpha$, meaning that j had to purchase the edge $\{j, c\}$. Similarly, when $k > 0$, then $\exists i \in N$ such that $c_i(s) = 2n - 3$, meaning that c covered the cost of the edge $\{c, i\}$.

Partition the set of agents N into three sets:

- $N_c = \{c\}$ is the singleton consisting of the central agent,
- $N_b = \{i : c \in s_i\} \subseteq N \setminus \{c\}$ is the set of agents who built an edge to c, and
- $N_{\neg b} = \{j : j \in s_c\} \subseteq N \setminus \{c\}$ is the set of "free-loading" agents who do not buy an edge to c.

Note that we have $|N_{\neg b}| = k$ and $|N_b| = n - k - 1$. All agents within a particular part are cost equivalent, so the cost to an agent i given the joint strategy profile s is

$$c_i(s) = \begin{cases} n - 1 + \alpha k & \text{if} \quad i \in N_c \\ 2n - 3 + \alpha & \text{if} \quad i \in N_b \\ 2n - 3 & \text{if} \quad i \in N_{\neg b} \end{cases}.$$

When $k > (n - 2)/\alpha + 1$, agents in N_c incur the highest cost and agents in $N_{\neg b}$ incur the lowest cost. Therefore, the inequality ratio is $(n - 1 + \alpha k)/(2n - 3)$. This ratio is maximized when $k = n - 1$ (i.e., in the center-sponsored star), so the maximal inequality ratio between agents in N_c and $N_{\neg b}$ is $(\alpha + 1)(n - 1)/(2n - 3)$, giving us

$$\lim_{n \to \infty} \text{IR} = \frac{1 + \alpha}{2}. \tag{3.5}$$

When $k < (n - 2)/\alpha$, the central agent in N_c incurs the lowest cost while agents in N_b incur the highest cost. In this case, the inequality ratio is $(2n - 3 + \alpha)/(n - 1 + \alpha k)$. This quantity is maximized when $k = 0$ (i.e., in the peripheral-sponsored star), so the maximal inequality ratio between agents in N_c and N_b is simply $(2n - 3 + \alpha)/(n - 1)$, giving us

$$\lim_{n \to \infty} \text{IR} = 2. \tag{3.6}$$

Hence, the largest inequality ratio for a star topology is the maximum between Equations (3.5) and (3.6) as $n \to \infty$. □

Fact 4.7. *A δ-regular graph of diameter two can have at most $n = \delta^2 + 1$ nodes.*

Proof. This fact follows from the Moore bound (cf. the survey [28]). For completeness, we summarize the relevant part of the presentation in [28, Section 3.1] regarding the Moore bound.

Consider a node v in a δ-regular graph G. Let n_i denote the number of nodes at distance i from v in G. We can bound $n_i \le (\delta - 1)n_{i-1}$. If G has a diameter D then

$$n_D = \sum_{i=0}^{D} n_i \le 1 + \delta + \delta(\delta - 1) + \cdots + \delta(\delta - 1)^D$$

$$= 1 + \delta(1 + (\delta - 1) + \cdots + (\delta - 1)^{D-1})$$

$$= \begin{cases} 1 + \delta \frac{(\delta - 1)^D - 1}{\delta - 2} & \text{if } \delta > 2 \\ 2D + 1 & \text{if } \delta = 2 \end{cases}.$$

Hence, when the diameter is $D = 2$, there can be at most $\delta^2 + 1$ nodes. □

REFERENCES

[1] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. "On Nash Equilibria for a Network Creation Game." *ACM Trans. Econ. Comput.* 2:1 (2014), 2:1–2:27.

[2] N. Alon, E. D. Demaine, M. T. Hajiaghayi, and T. Leighton. "Basic Network Creation Games." In *Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '10*, pp. 106–113. New York, NY, USA: ACM, 2010.

[3] E. Anshelevich, A. Dasgupta, J. Kleinberg, and Éva Tardos. "The Price of Stability for Network Design with Fair Cost Allocation." In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, 2004*, pp. 295–304. Los Alamitos, CA: IEEE Computer Society Press, 2004.

[4] R. Aumann and R. Myerson. "Endogenous Formation of Links Between Players and Coalitions: An Application of the Shapley Value." In *The Shapley Value: Essays in Honor of Lloyd S. Shapley*, edited by A. E. Roth. Cambridge, UK: Cambridge University Press, 1988.

[5] V. Bala and S. Goyal. "A Noncooperative Model of Network Formation." *Econometrica* 68:5 (2000), 1181–1229.

[6] J.-M. Baland and J.-P. Platteau. "Wealth Inequality and Efficiency in the Commons Part I: The Unregulated Case." *Oxford Economics Papers* 49:4 (1997), 451–482.

[7] J.-M. Baland and J.-P. Platteau. "Wealth Inequality and Efficiency in the Commons Part II: The Regulated Case." *Oxford Economics Papers* 50:1 (1998), 1–22.

[8] F. Bourguignon, F. H. Ferreira, and M. Walton. "Equity, Efficiency and Inequality Traps: A Research Agenda." *J Econ Inequal* 5:2 (2007), 235–256.

[9] S. Bowles. "The New Economics of Inequality and Redistribution." Federico Caffè Lectures. New York, NY, USA: Cambridge University Press, 2012.

[10] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. New York, NY, USA: Cambridge University Press, 1996.

[11] E. D. Demaine, M. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. "The Price of Anarchy in Network Creation Games." *ACM Trans. Algorithms* 8:2 (2012), 13:1–13:13.

[12] J.-Y. Duclos and A. Araar. *Poverty and Equity: Measurement, Policy and Estimation with DAD*. New York, NY, USA: Springer-Verlag and the International Development Research Centre, 2006.

[13] S. Ehsani, M. A. Fazli, A. Mehrabian, S. S. Sadeghabad, M. Safari, M. Saghafian, and S. S. Fadaee. "On a Bounded Budget Network Creation Game." In *Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '11)*, pp. 207–214. New York, NY, USA. ACM, 2011.

[14] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. "On a Network Creation Game." In *PODC '03: Proceedings of the 22nd Annual Symposium on Principles of Distributed Computing*, pp. 347–351. New York, NY, USA: ACM, 2003.

[15] C. Gini. "Variabilità e mutabilità (variability and mutability)." Technical Report, Studi Economico-Giuridici della R. Universita de Cagliari, 1912.

[16] S. Goyal. *Connections: An Introduction to the Economic of Networks*. Princeton, NJ: Princeton University Press, 2007.

[17] S. Goyal and S. Joshi. "Unequal Connections." *Int. J. Game Theory* 34:3 (2006), 319–349.

[18] M. O. Jackson. "The Stability and Efficiency of Economic and Social Networks." In *Networks and Groups: Models of Strategic Network Formation, Studies in Economic Design*, edited by B. Dutta and M. O. Jackson, Chapter 6, pp. 99–140. Berlin, Heidelberg: Springer-Verlag, 2003.

[19] M. O. Jackson. *Social and Economic Networks*. Princeton, NJ: Princeton University Press, 2008.

[20] M. O. Jackson and A. Wolinsky. "A Strategic Model of Social and Economic Networks." *JET* 71:1 (1996), 44–74.

[21] K. Jain and M. Mahdian. "Cost Sharing." In *Algorithmic Game Theory*, edited by N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, Chapter 15. New York, NY, USA: Cambridge University Press, 2007.

[22] E. Koutsoupias and C. Papadimitriou. "Worst-Case Equilibria." In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science, STACS'99*, pp. 404–413. Berlin, Heidelberg: Springer-Verlag, 1999.

[23] E. Koutsoupias and C. Papadimitriou. "Worst-Case Equilibria." *Computer Science Review* 3:2 (2009), 65–69. Originally appeared in STACS 1999.

[24] N. Laoutaris, L. J. Poplawski, R. Rajaraman, R. Sundaram, and S.-H. Teng. "Bounded Budget Connection (BBC) Games or How to Make Friends and Influence People, on a Budget." In *Proceedings of the 27th ACM Symposium on Principles of Distributed Computing, PODC '08*, pp. 165–174, New York, NY, USA: ACM, 2008.

[25] A. Mamageishvili, M. Mihalák, and D. Müller. (2013). "Tree Nash Equilibria in the Network Creation Game." In *Proceedings of the Algorithms and Models for the Web Graph: 10th International Workshop, WAW 2013*, Vol. 8305 of Lecture Notes in Computer Science, pp. 118–129. Cambridge, MA, USA: Springer, 2013.

[26] M. Mihalák and J. C. Schlegel. "Asymmetric Swap-Equilibrium: A Unifying Equilibrium Concept for Network Creation Games." In *Mathematical Foundations of Computer Science 2012*, Vol. 7464 of Lecture Notes in Computer Science, pp. 693–704, Berlin, Heidelberg: Springer, 2012.

[27] M. Mihalák and J. C. Schlegel. "The Price of Anarchy in Network Creation Games Is (Mostly) Constant." *Theory of Computing Systems* 53:1 (2013), 53–72.

[28] M. Miller and J. Širáň. "Moore Graphs and Beyond: A Survey of the Degree/Diameter Problem." *The Electronic Journal of Combinatorics* 20:2 (2013), 92.

[29] T. Piketty. *Capital in the Twenty-First Century*. Translated from French by Arthur Goldhammer. Cambridge, MA, USA: Belknap Press, 2014.

[30] T. Roughgarden. "How Unfair Is Optimal Routing?" In *Proceedings of the 13th annual ACM-SIAM Symposium on Discrete Algorithms,SODA '02*, pp. 203–204, Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2002.

[31] J. E. Stiglitz. *The Price of Inequality: How Today's Divided Society Endangers Our Future*. New York, NY, USA: W. W. Norton & Company, Inc., 2012.

[32] É. Tardos and T. Wexler. (2007). "Network Formation Games and the Potential Function Method." In *Algorithmic Game Theory*, edited by N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, Chapter 19, pp. 487–516. New York, NY, USA: Cambridge University Press, 2007.

[33] H. R. Varian. "Equity, Envy, and Efficiency." *J. of Econ. Theory* 9:1 (1974), 63–91.